

Self-Resonance Frequency
of
Single-Layer Air Core Solenoids

Open Source Hardware Engineering

Winter, 2020

Contents

1	Introduction	2
2	Simulation Details	8
3	Process Overview	11
4	Process Details	18
5	Accuracy	26
6	Useful Solenoids	30
A	Mathematical Symbols	33
B	Functions and Coefficients	35
C	Dielectric Effects	38
D	Bibliography	42

Chapter 1

Introduction

Nearly 500 electro-magnetic (E-M) simulations of air-core solenoids were used to study their first self-resonance frequency (SRF) as a function of solenoid geometry.

- Formulas have been developed which can predict simulated SRFs with a typical accuracy of $\pm 2\%$.
- Data from this effort compares favorably with published SRF measurements of real solenoids.
- Some conclusions reached in previously published work are found to be a result of sparse sampling of solenoid parameters. This study explored a broader sample of the parameter space.

Motivation

This article describes a study of the SRF of wire-wound, air-core solenoids (aka coils). Although there are many self resonances in a solenoid, the focus here is with the first or lowest in frequency.

In many cases, knowledge of SRF for a specific coil is necessary so that it can be avoided. Conversely, understanding the dependence of SRF on coil geometry is helpful in designing solenoids. Operating frequencies are often limited to some fraction of the SRF. From this perspective, the accuracy to which the SRF need be known is not high. Errors in the range of 20-30% might be entirely acceptable. That said, there has been a fair bit of research

into making more accurate predictions beginning in the late 19th century and continuing through today.

There are other reasons for knowing about SRF. For example, some published research concerns design of Tesla coil secondary and extra coils.

Scope

This study was limited to solenoids wound with bare wire in free space. The use of insulated wire (other than perhaps magnet wire), or coil forms alters the SRF. Only passing mention of these situations will be found in this article.

Attaching wires to solenoid terminals for connection to various circuits changes the SRF as well, and there is only limited discussion of this effect.

Close-Wound Coils Not Included

Coils with turn spacing less than 1.2 times the wire diameter were not simulated during this study. The SRF of such coils cannot be accurately estimated using the functions developed in this article. The range over which the results may be safely extrapolated below the 1.2 ratio value is also not known.

For The Pragmatist

This article may not satisfy the pragmatist. Accurately predicting SRFs of free-floating coils, disconnected from any other circuit components is the goal here. Connecting the coil to a circuit typically lowers the apparent SRF by 5-10%. The task of predicting in-circuit behavior is not taken on here in any detail.

Magnetic vs Electric vs Electromagnetic

One thing we learned from this study is that solenoid behavior at the first resonance is not solely an inductive (i.e. magnetic field) problem; nor is it a capacitive (i.e. electric field) problem. It is an electromagnetic field problem, and all of the terms in Maxwell's equations are in play. Specifically, electric fields can be induced by magnetic fields, and more importantly, vice versa.

$$\mathit{curl} \mathcal{E} = -\frac{\partial \mathcal{B}}{\partial t} \quad \text{and} \quad \mathit{curl} \mathcal{B} = \mu\epsilon \frac{\partial \mathcal{E}}{\partial t} + \mu \mathcal{J}$$

At low frequencies, magnetic fields induced by varying electric fields (aka displacement currents) are relatively insignificant, and the equations may be approximated:

$$\mathit{curl} \mathcal{E} = -\frac{\partial \mathcal{B}}{\partial t} \quad \text{and} \quad \mathit{curl} \mathcal{B} \approx \mu \mathcal{J}$$

As frequency increases, so does the rate of change of the E-field, and the E-field magnitude grows in proportion to current in the coil as well. At some point, both of these changes work together to make the approximation invalid.

It is tempting to identify correlations between SRF and aspects of solenoid geometry as being either capacitive or inductive in nature. In particular, the former ignores the possibility that significant magnetic fields are being induced by large, quickly varying electric fields.

Avoidance is the Best Policy

This article intentionally avoids any attempt to identify electric and magnetic field effects individually at SRF. No opinions are presented or implied as to whether such ideas suggested in existing literature are valid or not. Simulation data is treated empirically and analogies to inductive or capacitive effects are purposely avoided.

Normalization

SRF is a function of solenoid geometry: wire diameter, winding pitch, coil diameter and turn count. With the exception of turn count, linearly scaling these values simultaneously by the same amount scales SRF (inversely) by the scale factor. This allows geometry to be normalized to a convenient reference.

Early research discovered the total length of wire, and the solenoid's aspect ratio (length compared to diameter) are the two most important determining factors. As a result, solenoid geometry is often normalized relative to the total length of wire wound to form the solenoid. This study is no different and uses the same procedure.

It is also known that other geometric properties have a smaller effect. The extensive simulations carried out for this article shed more light on the importance of these other parameters.

Only air core solenoids with bare wire are considered here. Taking account of permittivity of insulation and coil forms is beyond the scope of this article.

Measurement

A solenoid operating at the lowest SRF presents a high impedance at its terminals, along with large, reactive electric and magnetic fields. These fields (at significant amplitudes) occupy a much larger volume than the coil itself. As a result, measurements are sensitive to the immediate environment – nearby objects can change the measured SRF significantly by altering electric and/or magnetic fields. Even the routing of wires used to connect test equipment has an impact. This leads to two questions.

- Can the measurement environment for SRF be standardized?
- How does one predict SRF when a solenoid is used in an application different from some standardized environment?

As early as 1902 the first question has been dealt with by measuring SRF with the coil disconnected and with the immediate environment empty to various degrees. The second question seems to have received little attention in research which is not surprising; there are a huge number of variables in the way a solenoid's immediate environment might be configured in an actual application.

In prior research, isolated solenoids have been measured by application of various configurations of electro-magnetic fields. At resonance, the excitation induces currents in the solenoid which result in large electric and magnetic fields in the vicinity of the test coil. The enhanced fields are then detected by probes or antennas, resonance being assigned to the frequency at which the probe/antenna produces the largest output.

Drawings and photos of various test setups appear to place the test solenoid in the near-field area of a source, typically another solenoid. The test probe (e.g. a small dipole or gas discharge tube) is also placed in the near field area of the coil. Some research states that the source and probe apparatus is located far enough away from the coil as to have little effect on

its SRF, but test data to support these claims is not given. A few simulations run as part of this study suggest these claims are likely valid.

Terminology

In this article, some coil geometry parameters are referred to by names which are different than in prior research.

L/D becomes Aspect Ratio

Solenoid length divided by diameter is referred to as *solenoid aspect ratio* or just *aspect ratio* and symbolized by A_s .

p/d becomes Winding Aspect Ratio

Winding pitch divided by wire diameter is called *winding aspect ratio*, symbolized by A_w . This term is never referred to as simply the aspect ratio; that's reserved for A_s .

Velocity Factor becomes Normalized SRF

Much of existing literature characterizes SRF relative to the frequency where total solenoid wire length is one-half wavelength. SRF is estimated by multiplying this frequency or wavelength by a correction factor. Drude came up with such a term – his function $f()$.

Others have associated this correction with the effective velocity of an electromagnetic wave traveling along the wound wire helix. As a result, the correction factor is sometimes referred to as a “helical velocity factor”, “velocity factor” or similar. This factor is applied to the SRF instead of the self resonant wavelength. Drude’s $f()$ is therefore equivalent to the multiplicative inverse of these velocity factors.

This article uses a name for the correction factor which is intended to be more generic and less tied to any particular theory: *normalized SRF*. The wavelength and frequency (in free space) associated with the coil wire being a half wavelength long in free space is

$$\lambda_w = 2l_w, f_w = \frac{c}{\lambda_w} = \frac{c}{2l_w},$$

and normalized SRF, f_n is defined as the ratio between actual SRF, f_r and the half-wavelength frequency f_w .

$$f_n \equiv \frac{f_r}{f_w} = \frac{2l_w f_r}{c}$$

This unit-less quantity is numerically identical to Knight's v_{hx} and equal to the inverse of Drude's function $f(n, A_s, A_w, \epsilon)$. Given normalized SRF and coil wire length, true SRF is given by

$$f_r = \frac{c f_n}{2l_w}$$

Matlab and Octave

A zip file is attached to the PDF file containing this document. In there are a set of data and script files for Matlab and Octave which implement the algorithms described in this article. Many PDF viewers provide the ability to extract the attached zip file, as does the `pdftk` utility.

Symbology

See the appendix for a complete listing of mathematical symbols used in this article.

Chapter 2

Simulation Details

Simulation models contained nothing but a solenoid in free space. Excitations and probes were implemented mathematically by the simulation software. No source or probe objects were present in the coil environment.

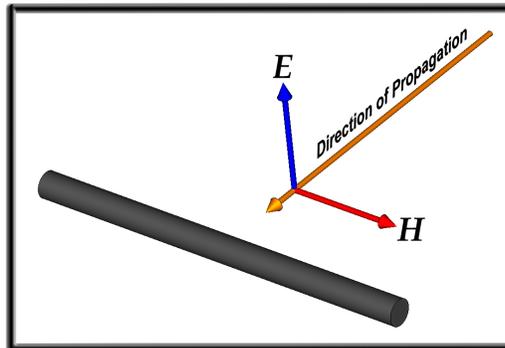


Figure 2.1: Simulated TEM plane wave orientation

Coil simulations were excited by an ideal TEM plane wave. At first, the plane wave was oriented as shown in figure 2.1. It was found however, that with very long coils (aspect ratios above ten or so) the E-field had to be rotated so that it had some component along the coil axis to get a significant response. Most of the later simulations had the E-field rotated by 45 degrees. Rotating the E-field angle did not appear to alter simulated SRF.

Resonance was defined as the frequency for which the H-field magnitude on the axis of the solenoid was highest (i.e. a local maximum).

Searching for SRF

One of the challenges with simulating free-floating coils was that the software wasn't really designed for this. It was not possible to get an interpolated frequency sweep in a single run which could be used to determine resonance. Instead, custom scripting was used which searched for the frequency at which the H-field was maximized. This technique required a separate simulation run at every test frequency, making simulations a time consuming process.

Meshing

Many of the early exploratory runs were performed with fixed tetrahedral mesh sizes. This allowed for faster runs but there was some question as to the accuracy of results.

Later, simulations to establish reference data used three passes of mesh refinement and compared favorably with the fixed-mesh runs.

Odd Behaviors

Geometries with very large wire diameters showed significantly different behaviors at resonance than more common designs. Specifically, most reasonable designs would show the peak in H-field strength along the axis was nearly coincident with a phase reversal of the H-field. For large wire diameters however, the phase reversal was at a substantially different frequency than the H-field peak. Also with large diameters, the Q of the coil seems significantly lower than for many other geometries. These two observations may be related.

Square Wire

Solenoids models used a square wire cross-section. Because simulations used a tetrahedral mesh, square wires could be more accurately modeled with a lot fewer mesh cells (compared to round wire). The resulting savings in simulation time was the motivation for this choice.

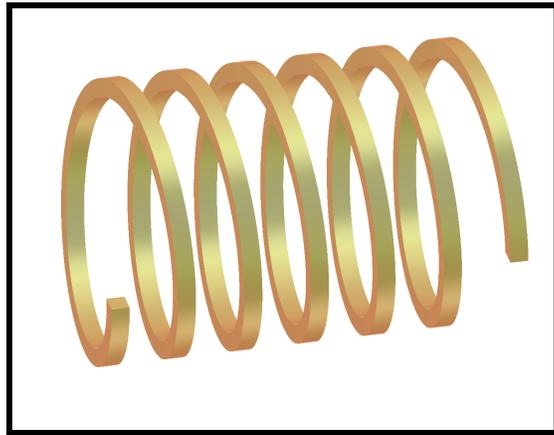


Figure 2.2: Simulated coil with square wire

Chapter 3

Process Overview

Here we present a general description of the process used to generate the estimation functions and some results. More detail on this is found in the following chapter.

Inside Diameter

Measurements of diameter and total wire length are based on the inside diameter of the solenoid. This is an empirically motivated choice – simulation data was easier to fit and estimate by using inside diameter. Although this may be correlated with skin and proximity effects, that connection is not suggested or relied upon here.

Sparse Sampling

Figure 3.1 shows the normalized SRF from about 500 simulations of free-floating solenoid, plotted against the solenoid aspect ratio (blue dots). A solid red line shows data from Drude (with winding pitch $A_w = 2.4$) and Pettit combined with smooth interpolation between those points. The Petit extension data for $A_s \geq 6$ is sparse and this results in a less than smooth plot there. The function developed by Knight to predict SRF is plotted as a solid green line.

For any given solenoid aspect ratio, there is a range of simulation data (blue dots) clustered around the value predicted by prior research. For aspect ratios above 8-10 or so, the range of simulated SRF increases significantly.

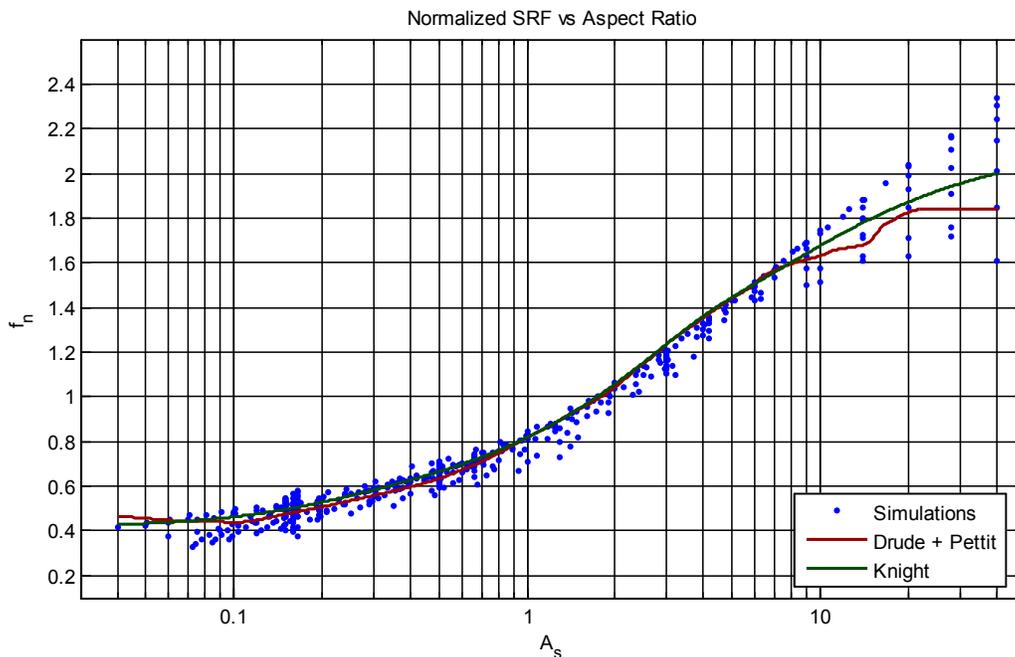


Figure 3.1: Summary of free-floating solenoid simulations

The key takeaway from this is that previously reported research represents a sparse sampling of the coil geometry parameter space. For example, Petit’s data appears to level off at high A_s values but also appears to be uneven or noisy. Simulations show that SRF at high aspect ratios (A_s) depends strongly on the winding pitch angle, and for coils with small pitch angles (high turn counts) normalized SRF keeps increasing and does not level off. It only begins to level off when the pitch angle is increased. Petit only measured coils with a single winding pitch angle for each large aspect ratio, and did not discover this dependence.

SRF Estimates

Three groupings of geometry parameters were found to be capable of accurately predicting normalized SRF found in simulations.

1. Solenoid aspect ratio, A_s
2. Winding pitch angle and solenoid aspect ratio taken together, (ψ, A_s)

3. Winding aspect ratio and turn count taken together: (A_w, n)

The normalized SRF estimate defined here is the product of three functions: an initial estimate P plus two corrections – Ψ corrects for winding angle, and Λ for the winding aspect ratio:

$$f_n = P(A_s) \Psi(\psi, A_s) \Lambda(A_w, n) \quad ,$$

where

P is one or more low order polynomials in $\ln(A_s)$.

Ψ is linear function of $\sin(\psi)$ where the slope and offset are determined by polynomials in A_s .

Λ is a non-linear function of both A_w and n .

If the two corrections are backed-out of the final estimate, we can write

$$P(A_s) = \frac{f_n}{\Psi(\psi, A_s) \Lambda(A_w, n)}$$

Similarly, the two corrections may be backed-out of simulation data like this:

$$\frac{f_{sim}}{\Psi(\psi, A_s) \Lambda(A_w, n)}$$

where f_{sim} is the normalized SRF found in the simulation. This result should ideally match the $P(A_s)$ function exactly. Put another way, backing out the two adjustments (Ψ, Λ) from simulation data allows it to be compared directly to $P(A_s)$.

The performance of SRF estimates developed in this article are depicted in figure 3.2. The same set of simulations shown in figure 3.1 (with corrections backed out) are plotted here against $P(A_s)$. This shows that SRF estimates accurate to 2% or better can be computed for most practical solenoid geometries.

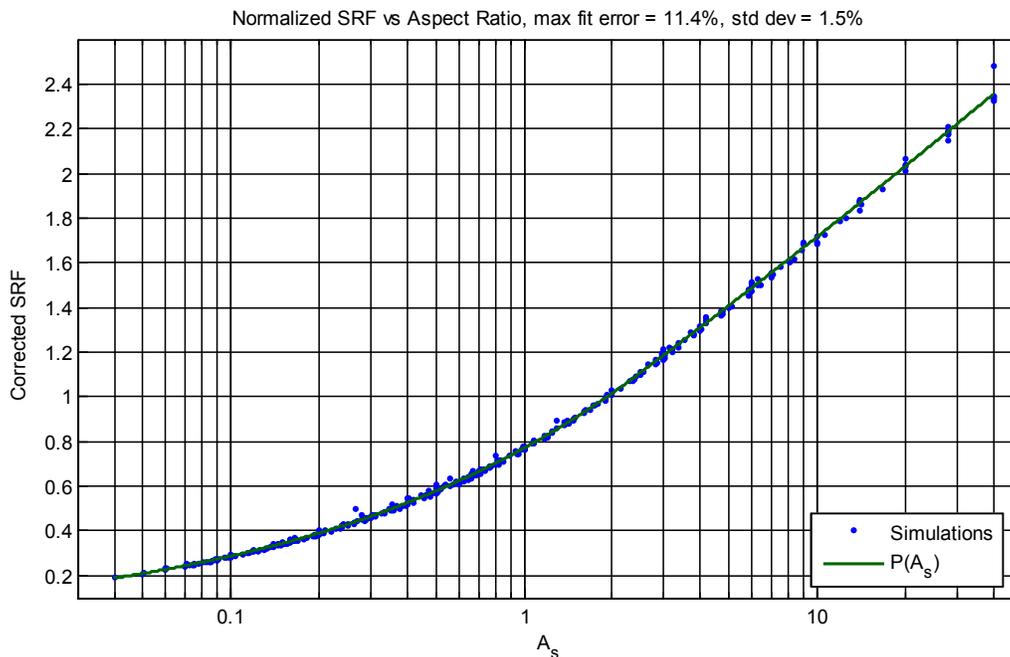


Figure 3.2: Simulations with $\Psi(\psi, A_s)$ and $\Lambda(A_w, n)$ backed out

Reference Points

A reference point is needed from which to develop the primary function P . One must specify where the correction functions Ψ and Λ are unity. The first correction is defined to be exactly one when the winding pitch angle is zero, regardless of the solenoid aspect ratio:

$$\Psi(0, \cdot) \equiv 1$$

The second correction is defined to be one for a winding aspect ratio of 2.4 and infinite turn count:

$$\Lambda(2.4, \infty) \equiv 1$$

The cases of course represent physically unrealizable coils. To approximate this ideal, the following guidelines were used for simulations to define an initial version of the reference polynomial $P(A_s)$.

- Winding pitch angle ψ preferably less than 0.02 radians but no more than 0.05 radians.
- A winding aspect ratio A_w of 2.4. For larger values of A_s this was not practical; when A_w was increased above 2.4, turn counts were kept high to minimize the effect.
- Turn counts preferably 40 or more. For very small solenoid aspect ratios, this was not practical. In those cases, turn counts were kept at 10 or more.
- Wire diameter should not exceed 2% of coil diameter. This was not suggested by the above, but used out of an abundance of caution.

Taken all together, this results in tightly wound coils using small diameter wire with high turn counts. At the extremes of solenoid aspect ratio, some of these goals must be compromised somewhat as indicated. Extreme values of A_s may not represent geometries that are useful in practice, so small deviations from the idealized references are perhaps less important.

One of the challenges occurs for large aspect ratios – it is not feasible to simulate the large turn counts that would be required for the target winding aspect ratio of 2.4:1. In these cases, A_w was allowed to increase to as much as 10:1.

Initial Estimate

The function $P(A_s)$ is defined in two overlapping piecewise segments with low-order polynomials. Splitting the range of aspect ratios into two segments avoids having to use a higher order polynomial to cover the entire range.

One segment covers aspect ratios below three, and the other is valid for all ratios greater than two. The polynomial argument is the natural logarithm of the aspect ratio:

$$P(x) = \sum_{j=0}^n c_j (\ln A_s)^j$$

Pitch Angle Correction: $\Psi(\psi, A_s)$

The first adjustment to the initial estimate of SRF is for the winding pitch angle, ψ . The correction is linear in either ψ or $\sin \psi$ ¹, while the coefficients themselves are polynomials in the logarithm of the solenoid aspect ratio.

$$\Psi(\psi) = c_1(A_s) \sin \psi + c_0(A_s)$$

Two fifth-order polynomials determine the linear slope c_1 , and offset c_0 as a function of aspect ratio.

$$c_0(A_s) = \sum_{j=0}^5 b_{0,j} (\ln A_s)^j \quad c_1(A_s) = \sum_{j=0}^5 b_{1,j} (\ln A_s)^j$$

Winding Aspect Ratio Correction: $\Lambda(A_w, n)$

This correction has significant dependence on both winding aspect ratio and turn count. For a number of different turn counts, the variation in SRF versus A_w has been approximated by cubic polynomials in the argument

$$Z(A_w) = \frac{1}{\cosh^{-1} A_w},$$

so for tabulated values of n , the correction is

$$\Lambda(A_w, n) = \sum_{k=0}^3 c_{n,k} Z^k$$

The first coefficient subscript in $c_{n,k}$ specifies the turn count, while the second specifies the power of Z in the polynomial.

Polynomial coefficients for several different values of n are tabulated in the appendix. When corrections are needed for a turn count other than a tabulated value, the following algorithm is recommended to interpolate the correction. Given a winding aspect ratio A and turn count m ,

1. Evaluate $\Lambda(A, n)$ for all tabulated turn counts. This yields several pairs of values $\{n, \Lambda(A, n)\}$.

¹It's not clear whether this is linear in ψ or $\sin \psi$ as pitch angles large enough for this to make much of a difference have not been simulated.

2. Fit a third-order polynomial to this set of data with n inverted: $\{n^{-1}, \Lambda(A, n)\}$.
3. Evaluate this polynomial with the argument m^{-1} to get the interpolated winding pitch correction.

Development Background

This correction was the most difficult to work out. Some of the thinking behind it is presented here. There exists a higher voltage and more phase difference between adjacent turns in coils with small turn counts. This might explain the strong dependence on turn count. This is only conjecture however, and no concrete mathematical model is offered to back it up.

Fitting the correction to values of Z is the result of guesswork; Z is the same variable in which the impedance of a parallel wire transmission line is linear. Lacking a more concrete theoretical underpinning, this should be considered a coincidence – a lucky guess. Perhaps a theory could be developed to explain this, but at this point it is just an empirical fact. If adjacent turns are truly acting as a parallel wire line, it is an odd one because each turn (except first and last) is adjacent to two other turns, and this is like no parallel wire line we know of. Furthermore, currents in adjacent turns are traveling in the same direction, whereas with a parallel line they would be in opposite directions.

Simulation scans to work out this correction held the solenoid and wire diameters, and turn count constant for each scan. The following limits were observed:

- $\psi \leq 0.02$
- $A_s < 1.0$
- $d/D = 0.01$

The first two items together are aimed at keeping the contribution of the Ψ correction constant and negligible. For each scan, the pitch was varied to accomplish the scan over values of A_w , in keeping with the above limits.

Chapter 4

Process Details

The process and data used to generate the three functions used in estimating SRF are detailed here.

$$f_n = P(A_w) \Psi(\psi, A_s) \Lambda(A_w, n)$$

Reference Scan

Reference scan data consists of pairs of $\{A_s, f_n\}$ values, covering a range of aspect ratios from 0.05 to 40. From this data, two piecewise polynomial fits to the raw reference scan data are created. This is a temporary fit which will eventually be discarded, and is named $Q(A_s)$ here. It is shown in figure 4.1.

The aspect ratio data is split into two overlapping ranges, $\{0..3\}$ and $\{2..50\}$. The polynomial argument is the natural logarithm of A_s . The first range requires a cubic polynomial, while the second range only needs a quadratic or linear fit. In figure 4.1 one polynomial fit is in red while the other is in blue.

The goal for this simulation scan was to keep winding aspect ratio $A_w = 2.4$ for the entire scan, but this would require very large turn counts for the higher values of A_s . As such, A_w was allowed to increase above 2.4 during the scan. In figure 4.1, a unique data point symbol/color has been assigned to each different value of A_w used in the scan.

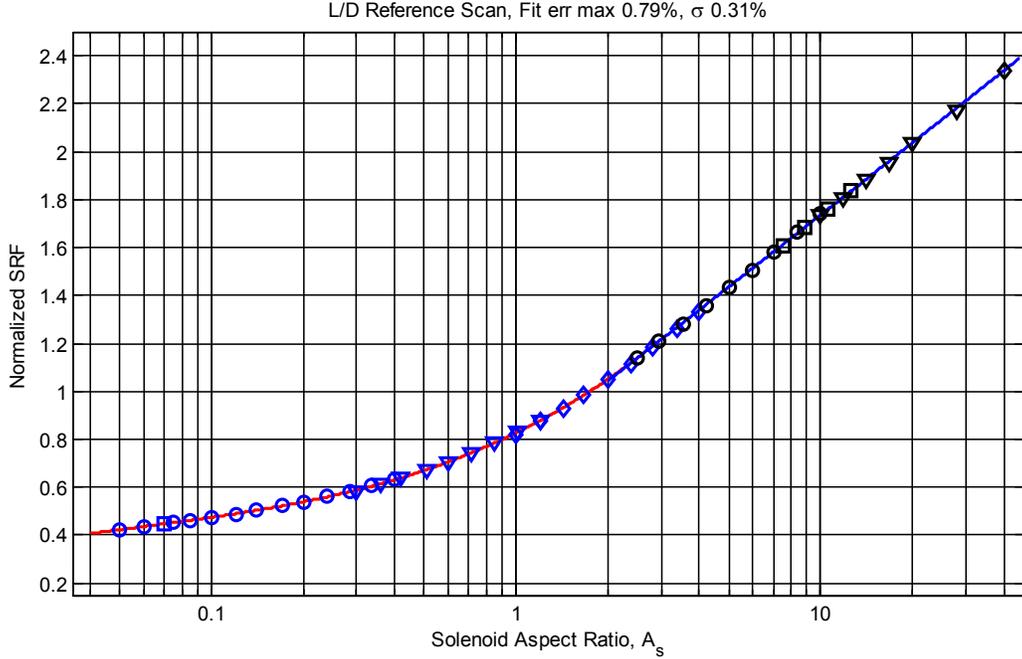


Figure 4.1: Polynomial fits to uncorrected reference scan

Generating $\Lambda(A_w, n)$

Data from scans of winding aspect ratio A_w was compared to the un-corrected polynomial reference fit $Q(A_s)$. The normalized SRF f_n was divided by the reference polynomial fit and that ratio is what was analyzed in this step: $f_n/Q(A_s)$. There are individual scans for different values of turn count n , where each scan contains simulations covering a range winding aspect ratios A_w . That data is graphed in figure 4.2, and covers a range of A_w from 1.2 to 7.0.

Cubic polynomials in $Z(A_w) = [\cosh^{-1}(A_w)]^{-1}$ were then fit to the data associated with each value of turn count. Eight different polynomial fits were produced. These are plotted as solid lines in figure 4.2.

Although the maximum value of A_w scanned was about 7.0, it is suggested (without justification) that the polynomial fit may be extended to A_w values as high as 100.

The smallest winding aspect ratio examined was 1.2, and we don't feel it is safe to extrapolate below this value. Many close-wound coils constructed

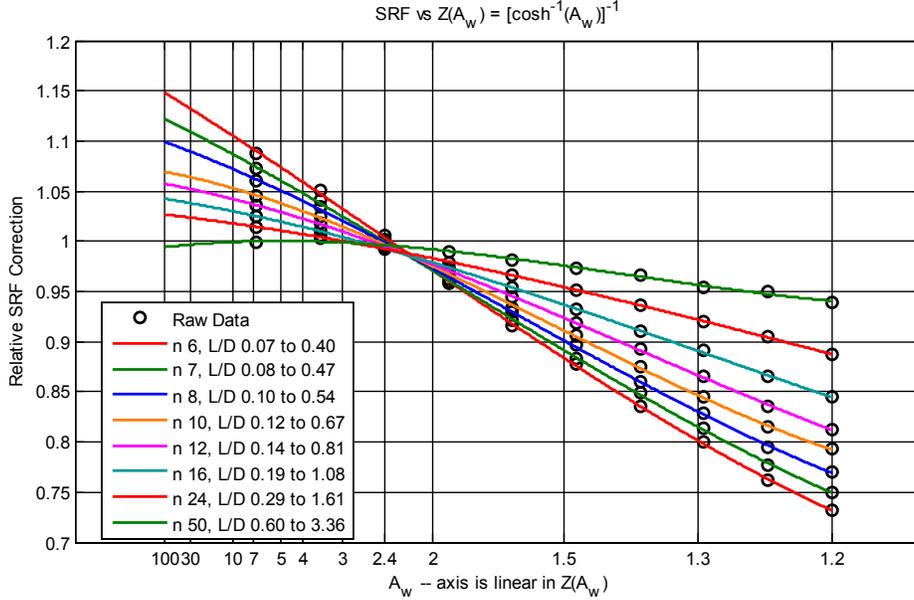


Figure 4.2: A_w corrections relative to reference scan

with magnet wire will have values of A_w below 1.2 and their SRF might not be accurately estimated using these methods.

If $\Lambda(A_w, n)$ is required for a tabulated turn count, the tabulated polynomial fit is simply evaluated. For non-tabulated turn counts, interpolation is required. The suggested process has two steps since the Λ function's domain is two-dimensional as shown in figure 4.3.

First, all polynomials are evaluated for the aspect ratio A_w of interest. This yields several (e.g. eight) pairs of data $\{n, \Lambda(A_w, n)\}$. Next, a second-order polynomial is fit to pairs of inverse turn count and correction values, $\{n^{-1}, \Lambda(A_w, n)\}$, and the polynomial is evaluated for the desired turn count.

Although the axes in figure 4.3 are labeled in turn count and winding aspect ratio, they are actually plotted linearly in Z and n^{-1} . This shows the function is smooth and well behaved in these variables, making interpolation a safe process.

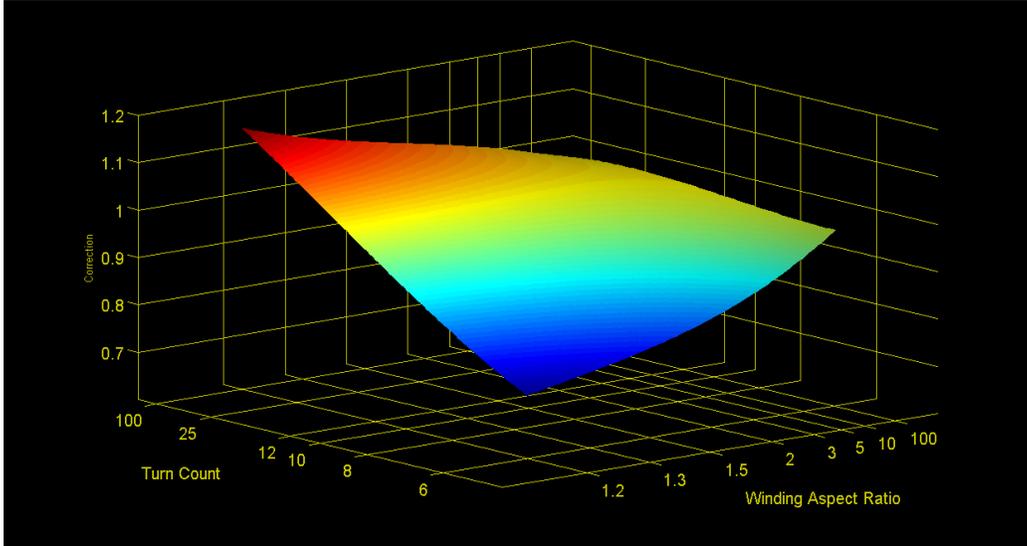


Figure 4.3: The Λ correction as a function of both arguments

Generating $\Psi(\psi, A_s)$

Scans of ψ were first corrected with the previously developed Λ function, then analyzed relative to the reference polynomial $Q(A_s)$. The data analyzed for this step is $f_n/[Q(A_s) \Lambda(A_w, n)]$. Scan data for this step is grouped by different solenoid aspect ratios, with each sub-scan covering a range of winding pitch angles.

This data is all nearly linear in $\sin \psi$ and can be approximated with a linear function. The slope and offset values for each value of A_s are plotted as a function of $(\ln A_w)$ in figure 4.5, along with cubic polynomial fits.

Since scans were only performed for aspect ratios of 3.0 and above, some anchor points have been added for small aspect ratios to prevent the curve fit from going off into the weeds there. For small aspect ratios, ψ is going to be very small and the actual slope represented by the anchor points is not all that important – it will be multiplied by a small value of ψ and won't contribute that much to the correction factor.

The offset used for anchor points is perhaps more important. There's not a lot of justification for the value actually chosen (1.06), other than the fact that net results of the Ψ correction are satisfactory. A value of 1.04 could have been used, but the polynomial fit is a little better with the anchors at

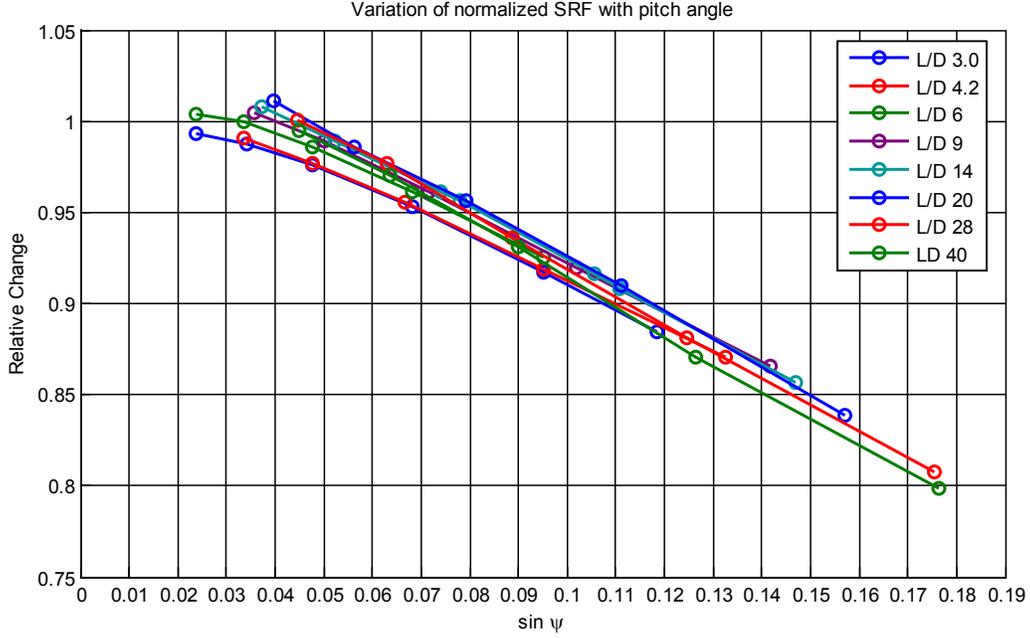


Figure 4.4: Ψ corrections relative to reference scan corrected for A_w

1.06. Finally, note that the variation in offset only covers a range of about 3% or so and this doesn't represent a very significant correction.

To compute $\Psi(\psi, A_s)$, the slope and offset polynomials are first evaluated with argument $(\ln A_s)$, yielding a slope, c_1 and offset c_0 . The resulting linear function is then evaluated using $\sin \psi$ as an argument.

$$\Psi(\psi, A_s) = c_1(A_s)\psi + c_0(A_s)$$

Generating $P(A_s)$

In this final step all simulations (about 500 in total) were used. The functions Ψ and Λ (as determined above) were backed out of the simulation data. The resulting data is then

$$\frac{f_n}{\Psi(\psi, A_s) \Lambda(A_w, n)}$$

That data presents a fairly smooth curve, with data points lying in a

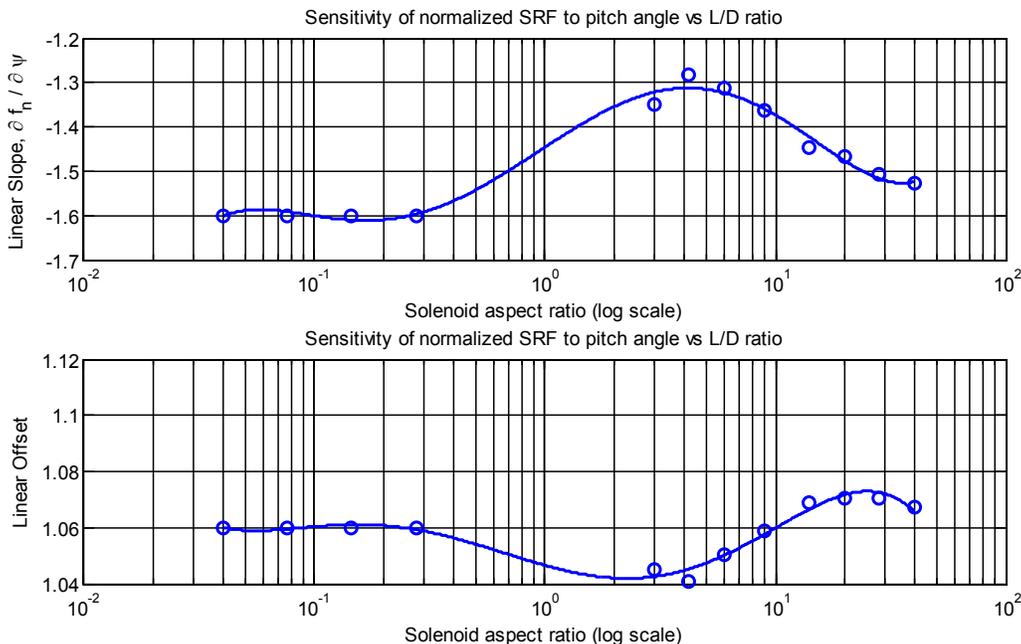


Figure 4.5: Ψ linear fit coefficients versus $\ln A_s$

narrow band at each value of A_s as shown in figure 4.6. Two polynomials covering the same overlapping ranges of A_s as used in $Q(A_s)$ are fit to this data, using $(\ln A_s)$ as the independent variable, to get the initial estimate function, $P(A_s)$.

It is important to note that the original polynomials $Q(A_s)$ fit to the reference scan (figure 4.1) are *not* exactly the same as the final $P(A_s)$ polynomials. $Q(A_s)$ is only a reference point from which Ψ and Λ corrections are computed. The two functions are almost the same because the goal of corrections was to bring data into agreement with $Q(A_s)$, but it's not exact – in particular because the Ψ correction has a value of about 1.06 (six percent high) for small values of aspect ratio A_s .

Figure 4.7 shows the difference between simulation data (corrected for Ψ, Λ) and the final $P(A_s)$ function.

This shows that the resulting composite SRF estimate has a typical accuracy in the range of 1-2%.

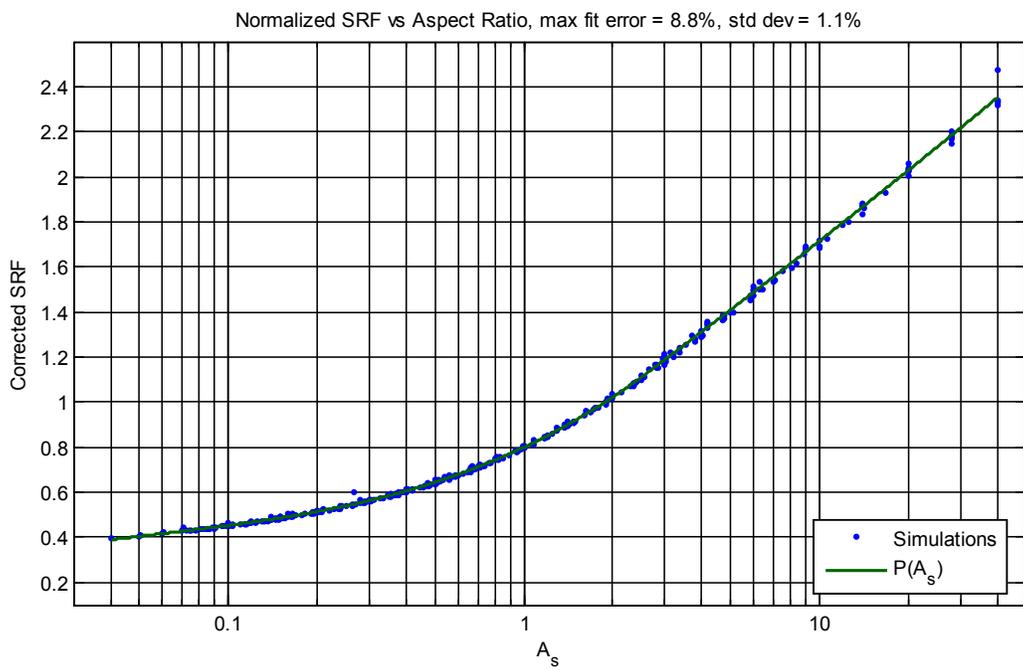


Figure 4.6: All simulation data with Ψ and Λ corrections

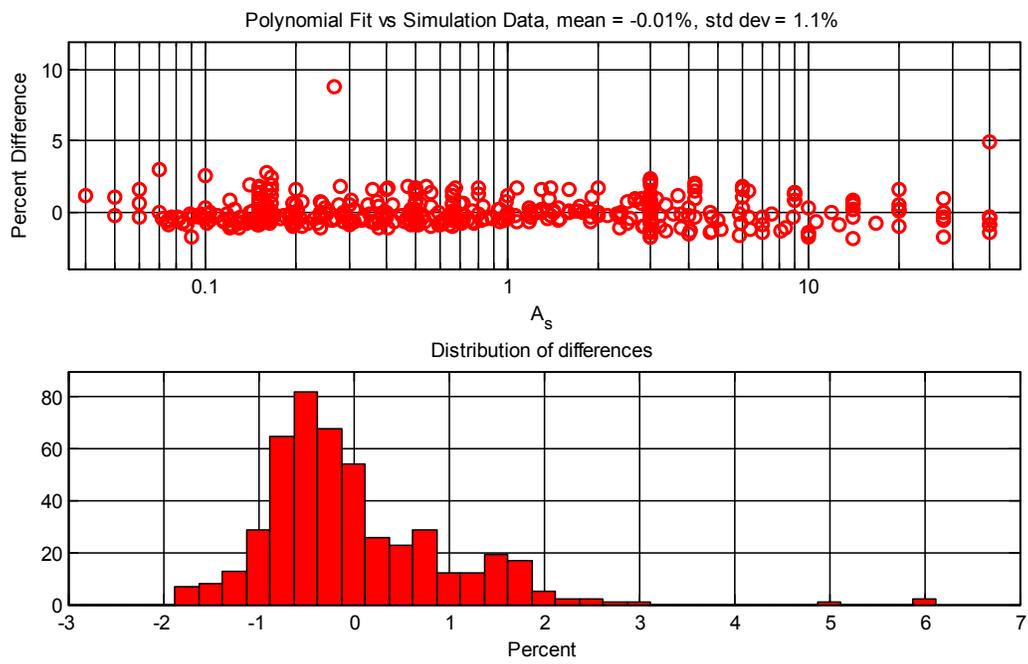


Figure 4.7: Final $P(A_s)$ fits vs simulation data corrected for Ψ , Λ

Chapter 5

Accuracy

While the corrections provided for winding pitch (Ψ) and winding aspect ratio (Λ) functions are significant, it is useful to compare the $P(A_s)$ function without these corrections to prior measurement data and the function developed by Knight.

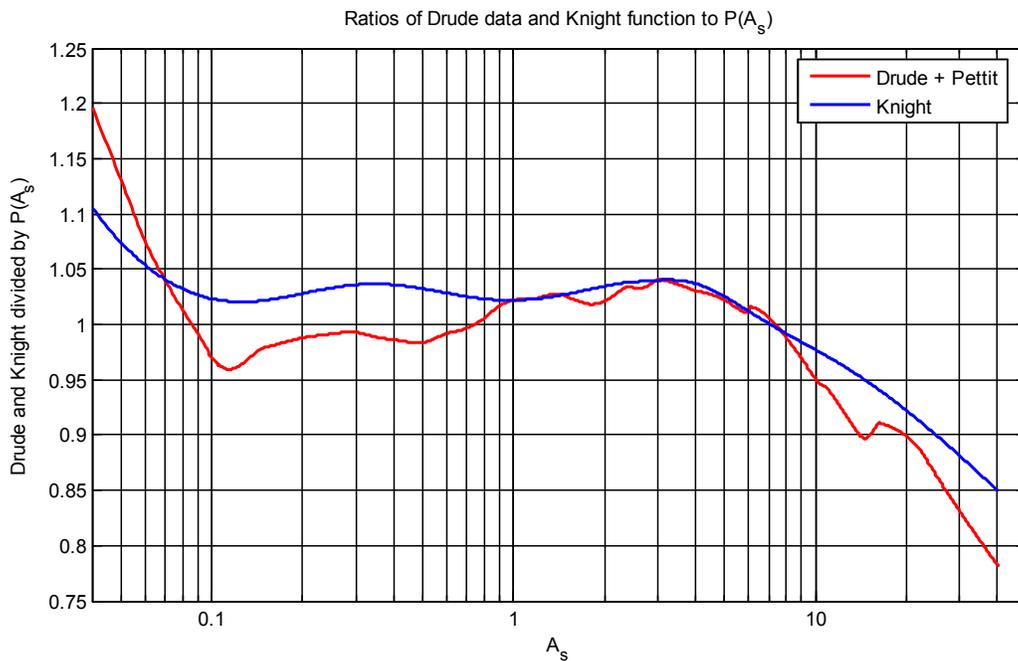


Figure 5.1: Knight and Drude normalized SRF relative to $P(A_s)$

Figure 5.1 shows the ratio between normalized SRF data from Drude and Knight's function, and the initial estimate function $P(A_s)$ developed in this study. There's reasonable agreement through the middle range of aspect ratios ($0.1 < A_s < 6$).

Drude's data begins disagreeing below about $A_s = 0.06$ for unknown reasons. Knight has commented that the flattening of Drude's curves at very low aspect ratios may be suspect. Also, at high aspect ratios, Petit's data ($A_s > 6$) likewise begins to diverge from $P(A_s)$. It has previously been noted this is likely due to a sparse sampling of coil geometry values. It may also be related to some extremely large pitch angles used by Petit at aspect ratios above 20.

Knight's function diverges at very low A_s values, but not as much as Drude's data and this is probably because Knight dismissed the flattening of Drude's data there. The Knight function also diverges at high aspect ratios, probably because it is based on Petit's data which has been shown to represent a sparse sampling of coil geometries.

Validation

Using numbers published by Knight for his 18T coil, and adjusting diameter and wire lengths to reflect inside diameter, we get an estimated SRF of 26.50MHz which differs from the published SRF measurement (26.69MHz) by -0.7%. There's no correction for round wire in this estimate, and increasing A_w from 1.787 to 1.912 makes the estimate exact. Being only a single data point, attempting to deduce a general round wire correction from this is not justified. Without more analysis, it's not even possible to conclude the difference is due to the wire profile.

Unfortunately, Drude did not publish turn counts for his test coils. As a result, it is not possible to compare his data with the estimates developed here.

Petit's Data

We also compared Petit's data to the uncorrected P function, by backing out the Ψ and Λ corrections from that data. That result is plotted along with the initial SRF estimate $P(A_s)$ in figure 5.2.

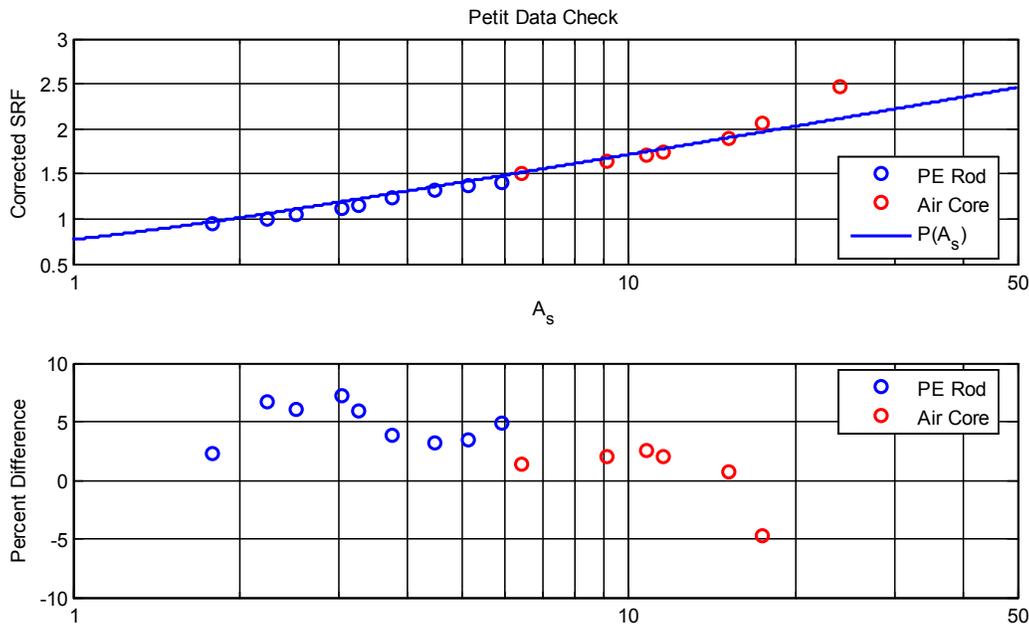


Figure 5.2: Comparing Petit's SRF data with estimates

Data taken with a PE rod core is shown in blue and the values have been increased by 5% to compensate for the presence of the polyethylene former rods. Air core data is red.

Agreement is reasonable for $A_s < 20$, but for $A_s = 24$ and 42 , the values differ from $P(A_s)$ by 15 and 41% respectively and are not visible on the plot. These represent very large pitch angles which we did not simulate when designing the estimating functions. There seem to be two possibilities; either more simulations are needed to reveal undiscovered behavior or some issues may exist with Petit's measurements. No attempt to resolve this issue is made as part of this study.

Dielectric Effects

See the appendix for information on a brief investigation that was done regarding coil forms.

Validation Summary

SRF estimates developed in this study appear to be in good agreement with measurements published by Knight and Petit.

Chapter 6

Useful Solenoids

It is encouraging to know that simulated SRFs of isolated solenoids can be predicted accurately, and that the simulations correlate nicely with measurements of physical solenoids. However, this is not worth much unless the results can be translated to the environment in which the solenoid is actually used.

A series of simulations were performed where the solenoid was driven by a discrete voltage source instead of being excited by an incident TEM plane wave. This necessitated adding wires to the simulation to connect the source.

As shown in figures 6.1 and 6.2, the feed wires were routed only a few wire radii above the coil wires themselves. Although this depicts round wires, square wires were used in the simulation runs.

The goal here was to examine SRF behavior in-circuit. This may or may not be a realistic model of actual use, but it serves to explore how SRF can change when a coil is not isolated in free space.

Resonance

One question that immediately arises when a coil is connected into a real-world circuit is: what defines self-resonance? One perhaps obvious answer is the frequency at which impedance is maximum – but at what point in the circuit is the impedance measured?

Particulars of the coil's connection into a circuit will alter the apparent resonance and accurately predicting that is beyond the scope of this article.

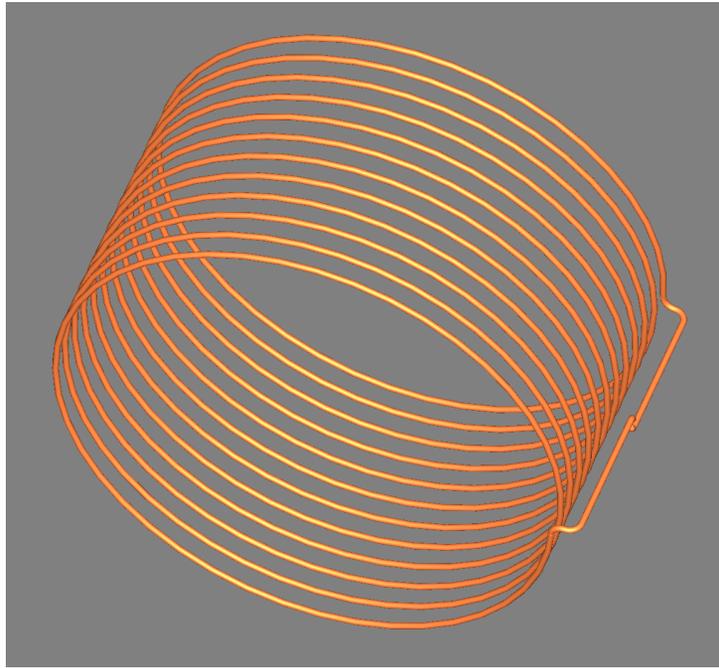


Figure 6.1: Typical model of driven solenoid

Source Connection Detail

Connection details for the discrete source port are shown in figure 6.2. The source is connected between the two tapered points visible in the center of the image.

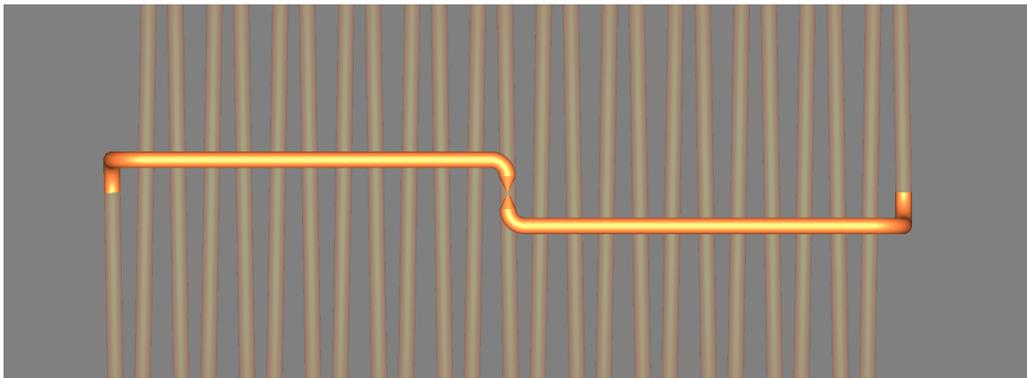


Figure 6.2: Discrete port connection detail

Using discrete sources in these models requires some care. Two physical points on the model are specified for the source connection. The source itself is modeled as having zero size, with infinitely thin, perfectly conducting wires connecting it to specified points in the structure.

In the limit, an infinitely thin wire has infinite inductance and this fact can cause instability in the simulation. It becomes necessary to prevent the software from increasing the mesh density around these wires. If the software starts increasing the mesh density in that area, the simulation becomes unstable and invalid results ensue. That is prevented in this case by placing the structure connections extremely close together, resulting in very short connecting wires.

As shown in figure 6.2, the feed wires are brought almost in contact with each other, using conical tips on the ends to provide very small landing areas for the discrete port connections.

It would have been possible to connect a short parallel-wire transmission line to the coil, and feed the line with a waveguide port in the simulation. Though this is in some aspects more realistic, electromagnetic fields at the end of the feed line would interact with the coil, probably generating different results. There would also be more configurable parameters such as the line's size and impedance to further confuse the issue.

Simulation Results

These runs were even more time consuming than with free-floating solenoids due to the extra connecting wires added.

Initial results seem to indicate that connected solenoid SRF values are five to ten percent lower than a free-floating coil. Not much work has been done with simulations here.

Appendix A

Mathematical Symbols

L Total length of coil winding, measured at center of wire.

D Coil (inside) diameter. Measured between inside faces of the wire, instead of center-to-center.

p Winding pitch. Center-to-center distance between adjacent turns.

ψ Winding pitch angle in radians, $\psi = \tan^{-1}(\frac{p}{\pi D})$

C Winding circumference = $\sqrt{p^2 + (\pi D)^2}$, based on inside diameter.

d Wire diameter, not including any insulation.

a Wire radius; half the diameter.

n Turn count.

l_w Total length of wire in the solenoid along inside diameter = nC

f Frequency in general.

f_r SRF in Hz.

f_n Normalized SRF (unitless).

c Without subscript, the speed of light in vacuum.

c_n With subscript, a general polynomial coefficient.

A_s Solenoid aspect ratio, L/D

A_w Winding aspect ratio, p/d

$\ln(x)$ is the natural logarithm of the argument x .

$\log_b(x)$ is the logarithm in base b of the argument x . There is no assumed base and this function is not used without specifying a base.

Appendix B

Functions and Coefficients

The SRF estimate is given by

$$f_r = \frac{c}{2l_w} P(A_s) \Psi(\psi, A_s) \Lambda(A_w, n)$$

Initial SRF Estimate

The $P(A_s)$ function is a piecewise polynomial defined as

$$P(A_s) = \sum_{k=0}^n c_k (\ln A_s)^k$$

Range	n	c_3	c_2	c_1	c_0
$A_s < 3$	3	8.16222×10^{-3}	$+ 7.00523 \times 10^{-2}$	2.68222×10^{-1}	7.98543×10^{-1}
$A_s > 2$	2		8.56326×10^{-3}	4.07654×10^{-1}	7.29018×10^{-1}

Figure B.1: Coefficients of $P()$ Polynomial

The two ranges overlap and either set of coefficients may be used in the overlapping region.

Winding Angle Correction

This correction is linear in $\sin \psi$, but the linear coefficients themselves are functions of the solenoid aspect ratio A_s . Fifth-order polynomials in $(\ln A_s)$ are used to determine them.

$$\Psi(\psi, A_s) = c_1 \sin(\psi) + c_0$$

where

$$c_k = \sum_{j=0}^5 c_{kj} (\ln A_s)^j$$

j	$c_{0,j}$	$c_{1,j}$
5	-1.38014×10^{-4}	$+1.12806 \times 10^{-3}$
4	-1.20550×10^{-4}	-5.54647×10^{-4}
3	$+2.58178 \times 10^{-3}$	-2.44538×10^{-2}
2	$+3.16921 \times 10^{-3}$	-4.56589×10^{-3}
1	-9.94120×10^{-3}	$+1.47443 \times 10^{-1}$
0	+1.04661	-1.44665

Figure B.2: Coefficients for determining c_1 , c_0

Although the sine of the winding pitch is used here, pitch angles are small enough in all of our simulations that the difference between ψ and its sine are quite small and the function might actually depend on ψ rather than $\sin(\psi)$.

Winding Aspect Ratio Correction

Cubic polynomials are provided for eight different values of turn count. Each one takes the argument

$$Z(A_w) = \frac{1}{\cosh^{-1} A_w} \quad ,$$

and gives the correction for that specific turn count:

$$\Lambda(A_w, n) = \sum_{k=0}^3 c_k(n) Z^k(A_w)$$

Unless the turn count of interest matches one of the values tabulated in figure B.3, the correction must be interpolated. To do this, first evaluate all eight polynomials at $Z(A_w)$, yielding eight pairs of values $\{n, \Lambda(A_w, n)\}$. Then fit a polynomial to the value pairs $\{n^{-1}, \Lambda(A_w, n)\}$. A first order linear fit is probably sufficient, but a 2nd or 3rd order fit can be made if desired. The Matlab/Octave script attached to the PDF file containing this document generates a 3rd order polynomial fit. That would look like this for an n th order fit:

$$\widehat{\Lambda}(A_w, n) = \sum_{j=0}^n \alpha_j n^{-j}$$

where the fit coefficients are α_j .

n	c_3	c_2	c_1	c_0
6	$+4.12860 \times 10^{-2}$	-8.55083×10^{-2}	-2.61242×10^{-1}	+1.20110
7	$+4.48689 \times 10^{-2}$	-1.11770×10^{-1}	-1.93482×10^{-1}	+1.16278
8	$+5.25879 \times 10^{-2}$	-1.47986×10^{-1}	-1.20904×10^{-1}	+1.12744
10	$+7.05875 \times 10^{-2}$	-2.16060×10^{-1}	-1.37143×10^{-2}	+1.07940
12	$+3.81287 \times 10^{-2}$	-1.35709×10^{-1}	-4.09260×10^{-2}	+1.06993
16	$+2.23880 \times 10^{-2}$	-9.02617×10^{-2}	-4.28932×10^{-2}	+1.05379
24	$+9.87100 \times 10^{-3}$	-5.06944×10^{-2}	-3.59897×10^{-2}	+1.03541
50	$+4.22761 \times 10^{-2}$	-1.48470×10^{-1}	$+1.05395 \times 10^{-1}$	+0.979135

Figure B.3: Λ polynomial coefficients

Appendix C

Dielectric Effects

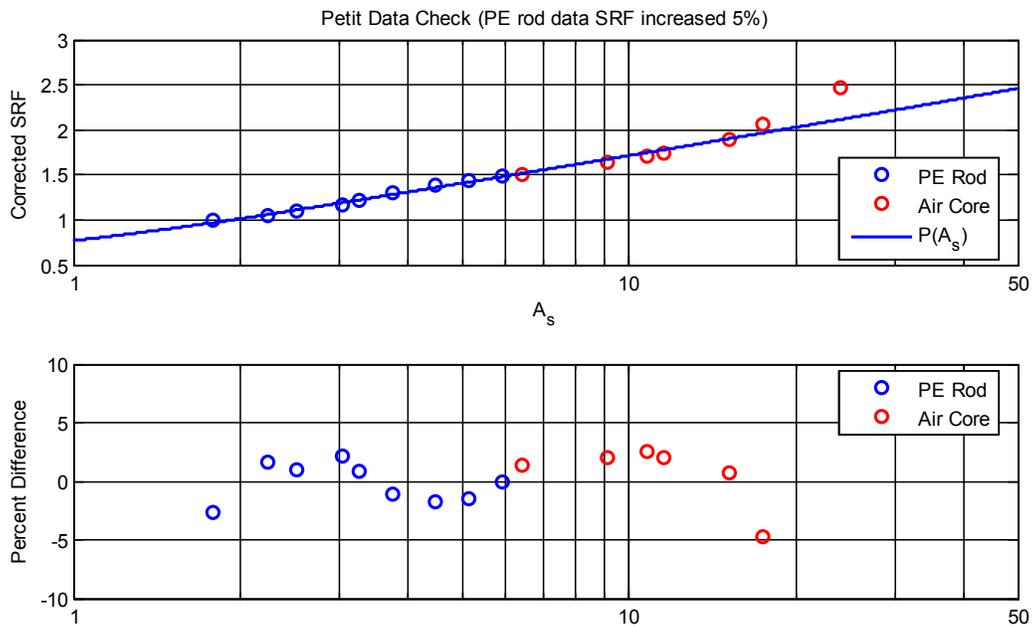


Figure C.1: Petit's PE rod data SRF bumped up by 5%

See the data shown in figure C.1. Here, Petit's SRF values which involved a coil former (PE rod) have been increased by 5%. The resulting excellent agreement with air-core data in the same plot suggests the addition of the PE rod lowered the SRF by 5%. However, there are no measurements of the same coil geometries without the rod, and all that can be said for certain is

that this 5% adjustment brings the data into agreement with data for large aspect ratios.

The Cadd Solenoid

Knight tested a Teflon-cored solenoid provided him by one Dr. Cadd. Buried in the blue data markers in figures 5.2 and C.1 is one point representing Knight's measurement of this Teflon-cored solenoid. Can you spot it? It's the middle point of the three which lie between $A_s = 3$ and $A_s = 4$. This data fits right in with Petit's – with no correction for the dielectric constant of Teflon. SRF estimated by the formulas developed here is 5.9% higher than what Knight measured (19.00 versus 17.94MHz).

While this study is largely concerned with air core solenoids, a few simulation runs were performed to investigate this further. Three runs which physically modeled the solenoid supplied to Knight by Dr. Cadd. Coil former dielectric constant was set to 1.0 (air or vacuum), 2.1 (Teflon) and 2.8 (ebonite) in three separate simulations. The wire was modeled in a groove one quarter of the wire diameter deep (1mm), with 44 turns instead of the published 43.8.

Figure C.2 shows the results of these simulations. For each line item, the difference between its SRF and SRF of other line items is shown – relative to the SRF of the line item. For example, differences on the air-core line item are relative to the air-core SRF. Column c for the air-core simulation line item indicates that SRF of the Teflon core simulation SRF is 4.8 percent lower – relative to the air-core SRF.

There are three main observations about this data.

1. Measured SRF was 1.2% lower than that predicted by the simulation with Teflon core. This suggests the simulations are reasonably accurate.
2. The Teflon core simulation shows a 5.0% increase in SRF when the core is removed.
3. The ebonite-core simulation showed a 7.9% increase in SRF when the ebonite core is removed.

A graph from page 93 of Knight is reproduced in figure C.3. Comparing the blue and green dotted lines, for the aspect ratio of the Cadd coil ($A_w =$

Item	SRF (MHz)	% Difference				
		a	b	c	d	
Physical Coil	a	17.939	–	-1.5	1.2	6.3
Ebonite-core Simulation	b	17.674	1.5	–	2.7	7.9
Teflon-core Simulation	c	18.160	-1.2	-2.7	–	5.0
Air-core Simulation	d	19.072	-5.9	-7.3	-4.8	–

Figure C.2: Simulations of the Cadd solenoid

3.1) we see about an 11% difference, while the simulations suggest the value is closer to 8%. This isn't conclusive, but looking at smaller aspect ratios, say $A_w = 0.4$, there is a 25% increase in SRF when the ebonite core is removed.

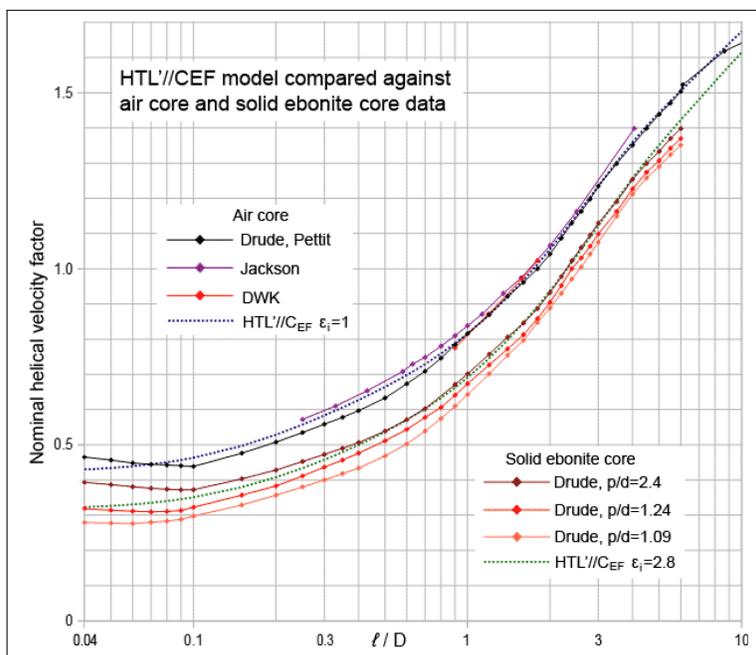


Figure C.3: From page 93 of Knight

Figure C.4 shows the results of these simulations. Two coil geometries were examined. The first is identical to the Cadd coil but with turn count reduced to only six turns. This gives an aspect ratio $A_s = 0.45$. Another two runs with and without ebonite core were performed with pitch and wire

diameter halved, and turn count doubled. This gave a different turn count but with identical values of solenoid and winding aspect ratios.

Item	SRF (MHz)	% Difference				
		a	b	c	d	
6T, ebonite core	a	55.899	–	19.5	–	–
6T air core	b	66.830	-16.4	–	–	–
12T ebonite core	c	29.298	–	–	–	24.4
12T air core	d	36.408	–	–	-19.5	–

Figure C.4: Modified simulations of the Cadd solenoid

In summary, simulations suggest that dielectric effects are roughly in line with that proposed in earlier research. There is some significant unexplained variability as shown in figure C.4 which is left unexplored by this study.

Appendix D

Bibliography

- [1] Drude, Paul Karl Ludwig, *On the construction of Tesla transformers, Period of oscillation and self-inductance of the coil.* (1902), Translation by Knight and Weaver, 2016.
- [2] Medhurst, R.G., *H.F. Resistance and Self-Capacitance of Single-Layer Solenoids*, *Wireless Engineer*, Feb 1947 pp35-43.
- [3] Knight, David W., *The self-resonance and self-capacitance of solenoid coils*, Self-published, DIO: 10.13140/RG.2.1.1472.0887