

Test of the Miller self-capacitance of a solenoid inductor

Bart H. McGuyer

<https://bartmcguyer.com/notes/note-12-MillerCapTest.pdf>

(Dated: January 31, 2021. Revision: June 1, 2021.)

TL;DR: Experimental demonstration of the “Miller” self-capacitance for a single-layer solenoid inductor in the lumped regime and a related result for weak loading. Involved modifying the solenoid to be a concentric cylindrical capacitor with variable capacitance and measuring its resonant frequencies with known capacitive loads.

This note presents the results of an experiment performed in 2014–2015 to test the “Miller self-capacitance” derived at least as early as 1919 by J. M. Miller^{1,2} and discussed in Ref. 3. This self capacitance is an analytical prediction for the effective capacitance of a single-layer solenoid that is grounded at one end and connected in parallel with a large capacitive load C_{load} . Importantly, its derivation assumes that the solenoid can be treated as an ideal transmission line with a constant, uniformly distributed capacitance c per length. Therefore, it is an approximation of the actual capacitance encountered in practice for which there is no accepted analytical model.⁴

The Miller self-capacitance has a particularly simple form. Its derivation³ assumes the voltage $V(x, t)$ along a solenoid of length H at any time is distributed linearly, $V(x, t) \approx (x/H)V_{\text{max}}(t)$. In a transmission-line model, the energy stored by capacitance is then

$$\frac{1}{2} \int_0^H c \langle V(x, t)^2 \rangle dx \approx \frac{1}{2} \int_0^H c \left(\frac{x}{H} \right)^2 \langle V_{\text{max}}(t)^2 \rangle dx = \frac{1}{2} C_{\text{Miller}} \langle V_{\text{max}}(t)^2 \rangle, \quad (1)$$

where brackets denote a time average. The Miller self-capacitance is

$$C_{\text{Miller}} = \frac{1}{3} C_0 \quad (2)$$

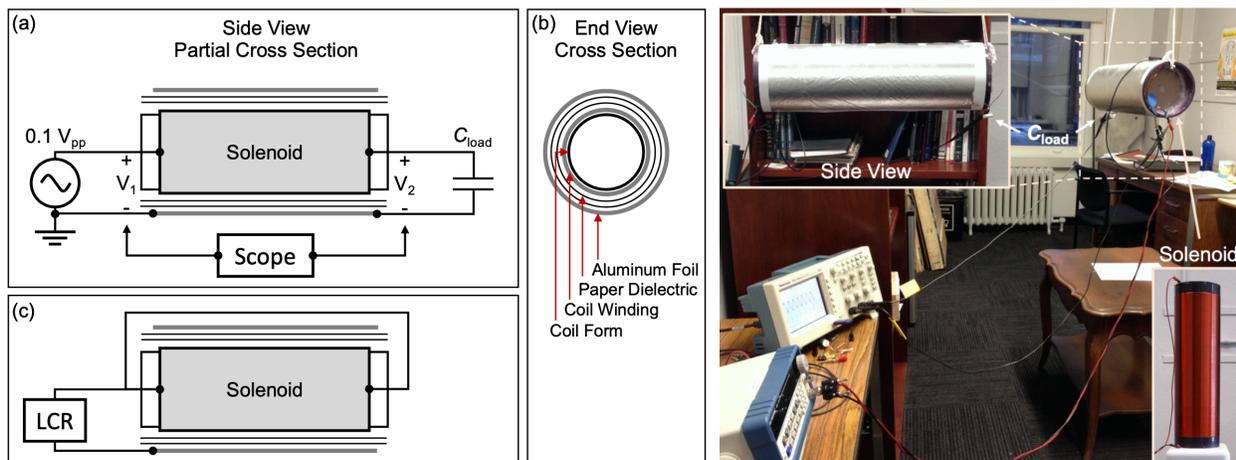


FIG. 1. Setup. (a) Arrangement to measure resonant frequency with C_{load} . (b) Cross section of solenoid showing outer layers of paper dielectric and foil to control the distributed capacitance c . (c) Arrangement to measure uniform capacitance C_0 . (Right) Picture of setup in arrangement (a).

expressed in terms of a uniform voltage capacitance for the solenoid

$$C_0 = cH. \quad (3)$$

The factor of $1/3$ comes from the integral $\int_0^1 x^2 dx = 1/3$.

I performed this experiment just out of curiosity to see if a real solenoid would reproduce the Miller self-capacitance, which is such a simple seeming result, in controlled conditions. Additionally, this experiment also probed a similarly simple result for the low-load limit that has a factor of $4/\pi^2$ instead of $1/3$.³

I. EXPERIMENTAL SETUP

Experimentally, the self capacitance of a single-layer solenoid is awkward to explore for a few reasons. First, at its simplest, it is a transmission line effect with the solenoid comprising one only of two wires of the line. The other wire is the solenoid’s environment, which can’t be ignored, in particular, because the solenoid’s ends are connected to ground and to C_{load} . So, the environment is an unspecified yet important systematic to control. Additionally, the self capacitance is usually in the pF range, which is . . . annoying to measure. Last but not least, you usually have to measure several solenoids to demonstrate any parametric dependence.

To address these issues, I chose to use only one solenoid and to vary its environment in a controlled way, as shown in Fig. 1. I did this by modifying the solenoid to be a concentric cylindrical capacitor by wrapping it with one or more layers of dielectric (almost always printer paper) and then aluminum foil as a grounded outer electrode.⁷ This approach brought both C_0 and the self capacitance into the nF range, which is much easier to accurately measure. To try to show the parametric dependence of (2) on C_0 with only one solenoid, I then varied the dielectric thickness and foil tension to control C_0 and measure its effect on the self capacitance of the solenoid.

Importantly, by significantly increasing the distributed capacitance c of the solenoid this way, the non-ideal effects of the inter-turn capacitances, primarily between neighboring turns, were artificially suppressed. Therefore, this setup should’ve made the solenoid behave a little more like an ideal transmission line, as assumed by the Miller self capacitance. However, no attempt was made to control the spatial variation of the distributed capacitance c along the coil, which is expected to peak near the ends, apart from keeping the solenoid reasonably far away from unwanted conductors.

Table I lists relevant physical parameters for the solenoid, which was made of a single layer of close-wound magnet wire on a PVC pipe form. Additional information about it is available in Ref. 5. Without a foil outer electrode, its self inductance L_s was independent of frequency to within measurement error from 100 to 100,000 Hz.

Measurements proceeded by first adjusting the outer dielectric and foil layers to get a desired value of the uniform capacitance C_0 , which was measured directly with an LCR meter (DER EE DE-5000) as shown in Fig. 1. Then the self capacitance was determined by measuring the lowest resonant frequency f_1 of the solenoid with each of three different capacitors of frequency-independent capacitance $C_{\text{load}} = 142.7 \pm 0.4$ (“150”), 48.6 ± 0.2

TABLE I. Physical parameters for the single-layer solenoid inductor.

Parameter	Value	Units	Comment
Diameter of winding	0.168 ± 0.002	m	Measured, pipe outer (winding inner) diameter.
Length of winding	0.506 ± 0.002	m	Measured.
Number of turns	712		Manufacturer label.
Wire gauge	22	AWG	Manufacturer label.
Resistance R_s	20.4 ± 0.2	Ohms	Measured, 100 Hz (DER EE DE-5000).
Inductance L_s	24.49 ± 0.09	mH	Measured, 10 kHz (DER EE DE-5000).

(“47”), or 5.02 ± 0.02 nF (“5 nF”). This used a signal generator (BK 4085) to drive the solenoid base with a 0.1 V peak-peak sine wave at several fixed frequencies about resonance, and an oscilloscope to measure the sinusoidal amplitudes of the base and top voltages of the solenoid with compensated 10X probes, as shown in Fig. 1. Preliminary tests found no difference between the results of this base-drive method and mutual-inductive or parallel-drive methods. This process was repeated for each value of C_0 . For $C_0 \approx 0.6$ nF, packing foam was used instead of printer paper. The lowest value of C_0 corresponds to a bare solenoid without dielectric or foil outer layers.

II. DATA ANALYSIS AND RESULTS

Fig. 2 shows example resonance data. To determine resonant frequencies, the frequency-dependent ratio of the sinusoidal amplitudes of the driving base voltage V_1 and the top voltage V_2 were fit with the expected resonant line shape for a driven series RLC circuit. Noting that the top voltage V_2 is across the total circuit capacitance C_1 , this line shape is

$$|V_2(f)/V_1(f)| = \frac{f_1^2}{\sqrt{(f_1^2 - f^2)^2 + (f\Gamma_1)^2}} \quad (4)$$

where $f_1 = 1/(2\pi\sqrt{L_s C_1})$ is an undamped resonant frequency and $\Gamma_1 = R_1/(2\pi L_s)$ is a damping factor from the total circuit resistance R_1 . Typical fit values of R_1 were around 150, 180, and 400 Ohms for 150, 47, and 5 nF C_{load} . As shown in Fig. 2, this line shape described the data well, but clear residuals similar to those shown were present in all data and increased in size with frequency.

The total circuit capacitance C_1 is assumed to be the sum of a self capacitance C_{self} for the solenoid and the load capacitance C_{load} ,

$$C_1 = C_{\text{self}} + C_{\text{load}}. \quad (5)$$

The values of C_{self} from each resonance were computed as

$$C_{\text{self}} = \frac{1}{(2\pi f_1)^2 L_s} - C_{\text{load}}. \quad (6)$$

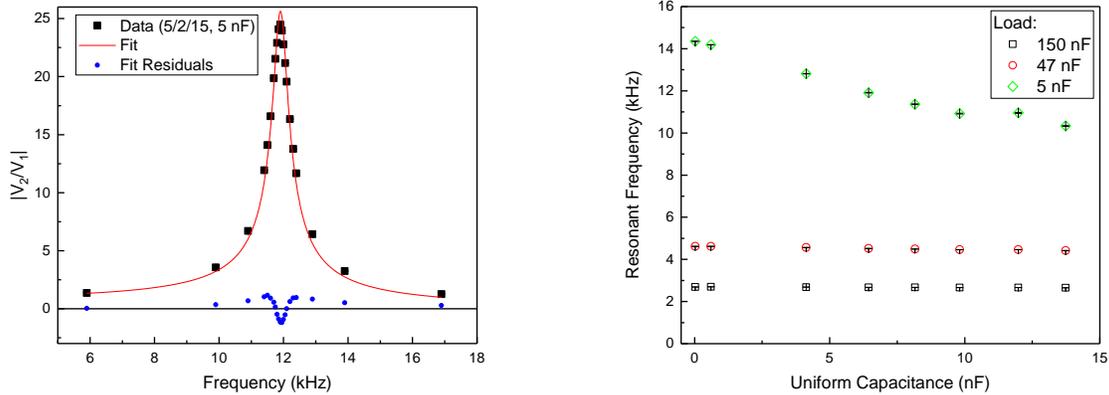


FIG. 2. Resonance data. (Left) Example fit to determine a resonant frequency f_1 , with fit residuals shown. (Right) All fitted resonant frequency data plotted against the uniform capacitance C_0 as directly measured by an LCR meter (at 1kHz above 2 nF, at 10 kHz below 2 nF).

The uncertainty for each value was computed by propagating the uncertainties from the solenoid self inductance, the resonant frequency, and the load capacitance as

$$\delta C_{\text{self}} = \sqrt{(\delta C_{\text{load}})^2 + (C_{\text{self}} + C_{\text{load}})^2 [(2\delta f_1/f_1)^2 + (\delta L_s/L_s)^2]}. \quad (7)$$

The load capacitance uncertainty tended to dominate. Note that one value of inductance was used for all data, so δL_s is a common error shared by all data points. Likewise, the uncertainty δC_{load} is shared by all data points using the same C_{load} .

It turns out that printer paper was a poor choice for the dielectric, because it has a significant dependence on frequency across the audio range. This is an important systematic error because the values of C_0 were measured at different frequencies than the corresponding f_1 and C_{self} . To correct for this error, I roughly characterized the type of printer paper used in an independent experiment. The results are summarized in Ref. 6 with a frequency-dependent capacitance $C(f)$ in terms of a Cole-Cole permittivity as

$$C(f) \propto \Re[\epsilon_r(f)] = \epsilon_0 + \epsilon_1 \frac{1 + \cos(\gamma\pi/2)(f/f_0)^\gamma}{1 + 2\cos(\gamma\pi/2)(f/f_0)^\gamma + (f/f_0)^{2\gamma}}. \quad (8)$$

For a capacitance normalized to unity at 100 Hz, $C(100 \text{ Hz}) = 1$, rough values for the parameters were $\epsilon_0 \approx 0.593$, $\epsilon_1 \approx 1.33$, $\gamma \approx 0.5$, and $f_0 \approx 24 \text{ Hz}$. Unfortunately, the frequency dependence of the actual dielectrics used couldn't be determined in situ because of the circuit behavior of the solenoid.

Fig. 3 shows the results. The top plot shows the values of C_{self} of (6) determined by fitting resonances versus the values of C_0 from LCR measurements. The bottom plot shows the same data, but the values of C_0 have been adjusted using (8) to their expected values at the corresponding f_1 for each data point.

III. DISCUSSION

Fig. 3 (bottom) presents the final results of this experiment. The data for each value of C_{load} tell a slightly story as discussed in detail below. Overall, the data do seem to provide support for the simple form of the Miller self-capacitance in (2) in such controlled conditions, and also to provide support for a related self capacitance near the unloaded limit.³

The 47 nF data (red circles) provide the most compelling evidence. For the full range of the data, the load was large enough to be in the lumped regime [following Fig. 13 (top right) of Ref. 3]. With the exception of a suspected outlier at $C_0 \approx 12$ nF (not adjusted) to be discussed next, the data agrees with the expected slope of $1/3$ to well within error.

The data at $C_0 \approx 12$ nF (not adjusted) appears to be an outlier. This is clearest in the top plot of Fig. 3, which has a dashed ellipse circling those data. Looking back over my original notes during the measurement, the main difference between that data and all others appears to be that I attempted to fine tune the value of C_0 by loosening the tension of the outer foil layer significantly. Unfortunately, I did not double check the value of C_0 after the resonant frequency measurements, as I usually did with other data points. Therefore, a potential explanation is that the outer layers were disturbed as the wiring to the solenoid was changed, altering the value of C_0 after it was measured and before the resonant frequency measurements. This would be consistent if C_0 decreased between measurements by about 3 nF ($\sim -25\%$), which is not small but not inconceivable.

The 150 nF data (black squares) have such large error bars that they appear nearly useless. However, the variation between data points seem to be much less than the error bars, again excluding an outlier at $C_0 \approx 12$ nF (not adjusted). Numerically, the error bars are mostly due to the uncertainty $\delta C_{\text{load}} \approx 0.003 C_{\text{load}}$, which is shared by each point. (Notice, for example, that the error bars are roughly uniform in size, even at the lowest value of C_0 .) This suggests that the error bars are instrumentally limited by the lack of knowledge of C_{load} , so they overestimate the measured slope uncertainty. The negative values of the data for $C_0 < 2$ nF and the overall appearance of the data being parallel to but offset from the dashed line provide some support for this.

Last but far from least interesting, the 5 nF data (green circles) have rather small error bars and deviate significantly from the Miller capacitance after C_0 is adjusted. However, note that once C_0 is near or above 5 nF (adjusted), the Miller capacitance is no longer the expected result for this data. Instead, as C_0 increases above C_{load} , the self capacitance in a transmission-line model should transition towards a slope of $4/\pi^2$, as shown in Fig. 13 (top right) of Ref. 3. That is what appears to occur, except again for the suspected outlier. Below about 5 nF (adjusted), the data slightly favor the Miller capacitance, but the error bars are comparable to the difference between both slopes. Out of curiosity, I did manually test the $(4/\pi^2)C_0$ result also by measuring the bare solenoid with no load ($C_{\text{load}} = 0$), but while it did seem to agree, I was not as confident in my pF-range error analysis.

Interestingly, the trend of higher C_{load} leading to lower self capacitance in Fig. 3(top) is mostly but not completely explained by the adjustment for the frequency-dependent dielectric in Fig. 3(bottom). Much of this may be due to a systematic offset from an instrumentally limited precision δC_{load} , as discussed above. Additionally, there might be a contribution from

error due to the resonance fit residuals, which may depend on frequency and thus also C_{load} .

Overall, the most significant challenge of this surprisingly tedious experiment was discovering and then correcting for the systematic error of using printer paper. Improved results should be possible by carefully choosing a more frequency-independent and lower loss dielectric to wrap the solenoid. Finally, one amusing phenomenon of the setup was that resonance could be heard audibly with large driving voltages ($\gtrsim 20 V_{\text{pp}}$).

ACKNOWLEDGEMENTS

I am grateful to Tanya Zelevinsky and Ken Sikes for their support with equipment and space, and to Mickey McDonald for helpful discussions.

REFERENCES

- ¹J. M. Miller, “Electrical oscillations in antennas and inductance coils,” *Proc. IRE* **7**, pp. 299–326 (1919). DOI: 10.1109/JRPROC.1919.217442
- ²V. G. Welsby, *The Theory and Design of Inductance Coils*, 2nd ed., Wiley (1960). See p. 150.
- ³B. H. McGuyer, “Deriving the equivalent circuit of a Tesla coil,” technical note, 2020.
Available online: <https://bartmcguyer.com/notes/note-11-TcEquations.pdf>
- ⁴T. H. Lee, *Planar Microwave Engineering*, Cambridge, (2004). See the footnote in Chapter 6 on p. 142.
- ⁵Bart McGuyer, “(Dataset) Measured self-resonant frequencies of an air-cored single-layer solenoid inductor,” Harvard Dataverse, V1 (2020). <https://doi.org/10.7910/DVN/QV8HQS>
- ⁶B. H. McGuyer, “Frequency-dependent capacitors using paper,” *Am. J. Phys.* **89**, 370 (2021).
Available online: http://bartmcguyer.com/pdf/2021_AJP_McGuyer_PaperCap.pdf
- ⁷I’m reasonably confident this outer electrode did not form a closed cylinder that might disturb the solenoid’s inductance. In initial tests, I used paper strips to keep overlapping foil from touching, but I didn’t take enough pictures to confirm this was done in the final measurements beyond all doubt.

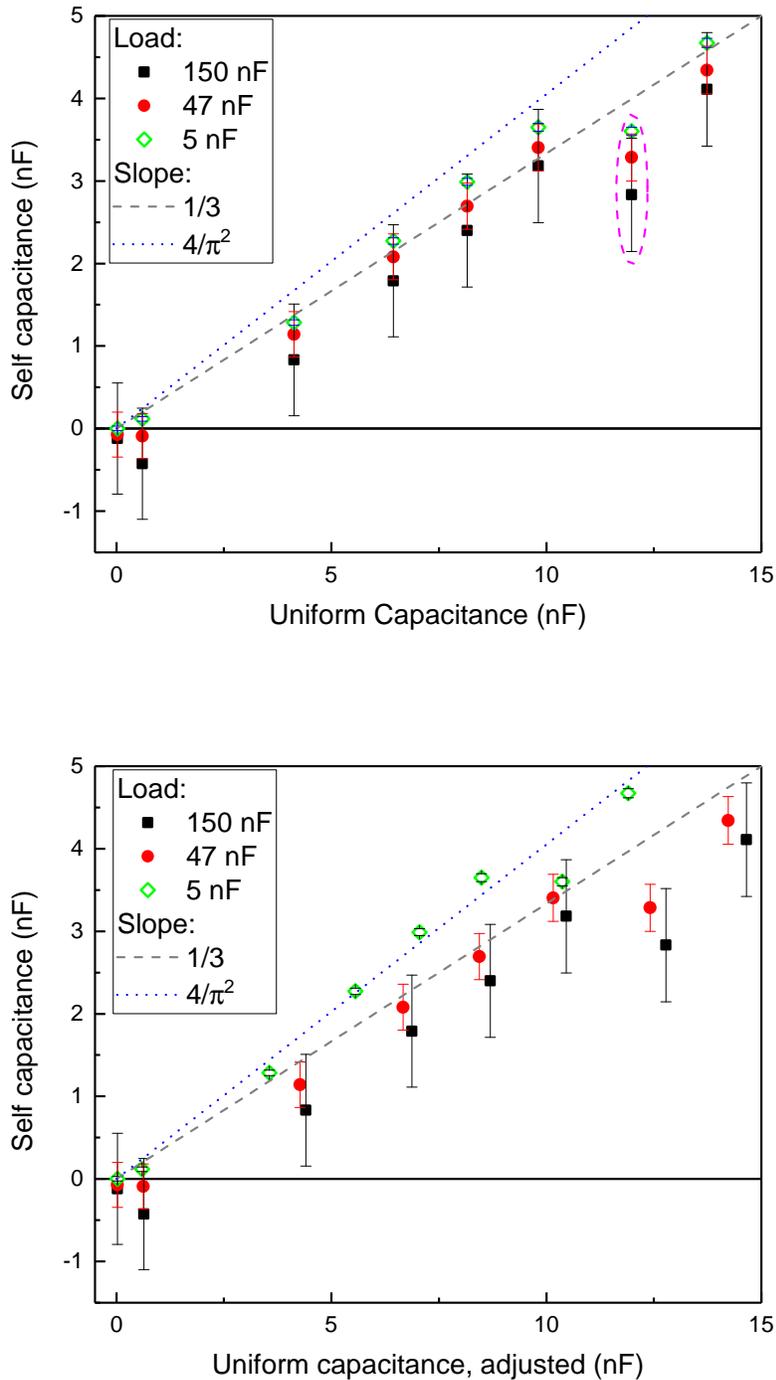


FIG. 3. Test results. Both plots show the solenoid self capacitance C_{self} versus the uniform capacitance C_0 . Dashed lines show the slope of $1/3$ expected for C_{Miller} in (2) and the slope of $4/\pi^2$ expected for an unloaded transmission-line model from Ref. 3 (c.f., Fig. 13 top right). The circled data for $C_0 \approx 12$ nF are suspected to be outliers from a disturbance to the foil electrode between C_0 and f_1 measurements. In the bottom plot, the values of C_0 have been adjusted to attempt to remove a systematic error from using a frequency-dependent paper dielectric.