Sanjaya Maniktala, 2013

Understanding and using LLC Converters to Great Advantage

PART 1: Overview of the Transition from Conventional PWM Power Conversion to Resonant Topologies

Introduction

At an early moment during the development of conventional "square-wave" switching power conversion, it was said that nature abhors any attempt to suddenly turn OFF the current passing through an inductor. That fear prompted early attempts at creating "resonant topologies". But for various reasons, in particular a very wide, almost uncontrolled switching frequency and resulting EMI spread, resonant topologies did not become mainstream, and *still aren't*. But nowadays there is a sudden resurgence of interest, primarily through the "LLC topology". This unique combination of two inductors and one capacitor ("L-L-C"), offers a relatively narrow range of switching frequencies, which are much easier to design a standard EMI filter for, combined with the capability of producing zero-voltage switching (soft-switching) through careful design which can significantly improve EMI and efficiency over a wide load range. But the overall topic and analysis is still considered complex, almost mystical, and is poorly understood. A majority of successful LLC designs are simply based on trial and error, i.e. the art of tweaking, both in a real lab environment, and in a virtual lab using Spice-based simulators. This has significantly hindered their adoption in a modern design-and-development, flowchart-based, commercial product environment.

To tackle the underlying "fear of resonance", in these pages, emphasis is laid on first understanding the guiding principles of resonance and soft-switching. Only after acquiring the underlying intuition and analytical depth, do we start to build the LLC converter, and we do that in steps: assembling the building-blocks one by one, at each stage analyzing the performance, providing the necessary simplified math, and validating each step via Mathcad and Simplis (Spice). In the process, a very unique method is being proposed on these pages perhaps for the very first time.

Soft and Hard Switching

To close loose ends, we need to revisit some of the originally expressed anxiety against interrupting inductor current in "square-wave" switching power conversion. That fear too is admittedly somewhat true, because energy is stored in an inductor as $\frac{1}{2} \times L \times I^2$, which depends on the current passing through it at any given moment. Since by the Law of Conservation of Energy, energy cannot just be "willed away" in an instant, we should never try to interrupt the inductor current by simply placing a switch (or FET) in series with it and turning that OFF. That attempt *will* destroy the switch --- on account of a huge voltage spike (induced voltage) which will suddenly appear. But there is a way out as people soon discovered: *we can consciously provide an alternate "freewheeling" path for the inductor current --- to keep it flowing smoothly (without any discontinuity) once the switch turns OFF*. That led to the "catch diode" seen in conventional switching power conversion. The provision of this diversionary diode path, a detour in effect for inductor current, is a much wiser option than trying to contain the current in a cul-de-sac of sorts, leaving it no way out. And that diode is the reason we can, by now, manage to get away turning the switch (FET) of any conventional power converter, ON and OFF, repetitively, hundreds of thousands of times every second, and reliably so. Nature is clearly not complaining. But there *was a price to pay* for this inductor-driven brinksmanship of sorts. Diverting the current at the very last moment, and we mean *very last moment* here, does create *penalties*. Because whenever we try to turn OFF inductor current, by say dragging the Gate of the series FET to ground (or to its Source terminal as applicable), we do *not* manage to *instantaneously* reduce the current in the FET as we had perhaps hoped to do. On closer analysis we see that there needs to be a full rail-to-rail *voltage swing completed* (across the FET), just to be able to forward-bias the catch diode, to get it to even *start* taking up some of the current away from the FET. But before that happens, we cannot afford to somehow force the FET to relinquish its current, for fear of the induced voltage effect. In other words, there is a transitory moment where there is a full voltage swing taking place across the FET. Though the average value of this voltage swing is half its peak-to-peak value, this is present with the *full* pre-existing inductor current continuing to pass through the FET. There is *no* reduction in FET current *yet*, because the diode is still in the process of being "set up" to *start* taking up slack. And that does not even happen till the *entire* voltage swing is *completed* and the diode gets *forward-biased*. Geometrically speaking, we see an "overlap" of voltage and current across the FET as shown in Figure 1, and that constitutes "crossover" (switching) losses in the FET. Mathematically speaking, there is a non-zero V x I product across the FET during the switch transition. Yes, there are other subtle contributors to this crossover loss component too, such as the forcible discharging (burning up of energy) of the parasitic capacitance across the FET, and also the Drain to Gate "Miller effect" causing an increase in switching (transition) time and causing even more V-I crossover energy to be wasted, etc. Therefore, as we go to higher and higher switching frequencies, this useless "work" done by the input source (in the form of dissipation), leads to several percentage points of degradation in overall efficiency. It is for these reasons that the small, but extremely significant moment of switch transition, is garnering a lot of attention today, piquing renewed interest in resonant topologies which offer hope in this regard.

The reduction of switching losses during the few nanoseconds of switch state-transition (crossover of the FET), is critical to improving conversion efficiency. But we are also hoping to reduce EMI by the softer "resonant transitions", since we know that conventional "hard transitions" are the main source of most of the EMI in conventional converters. We are almost intuitively expecting that by using resonant topologies, the voltage will be softly reduced by self-resonant action, so that the V-I "overlap" will be likely reduced, or become *almost zero*, as also displayed in Figure 1. If so, that should help reduce EMI too.

Note: Another way out, suitable for low power applications, is to use discontinuous conduction mode (DCM), because we can thereby ensure the current is zero whenever we turn the FET ON. We can implement this using the well-known "ringing choke converter" (RCC) principle, which basically senses the ringing voltage of the inductor/transformer, to gauge when there is no residual current left (i.e. core is de-energized), and turns the FET ON at that moment. This is variable frequency PWM, in the form of boundary-conduction mode (BCM). It was actually used on a very wide scale in the 1980s in the form of the ubiquitous and historic Television Power Supply flyback controller IC, the TDA4601, originally from Philips, and still available from Infineon. The modern iPhone charger uses the L6565 from ST microelectronics, which though based on the same old ringing choke principle of the TDA4601, prefers to call itself a QR (quasi-resonant), ZVS (zero voltage switching) topology SMPS controller chip, in keeping with the times. But there are switching losses when we turn OFF the FET!

Understanding and using LLC Converters to Great Advantage

Keep in mind however, that the resonant soft-switching sketched intuitively in Figure 1, is only a “wish list” so far. It is all much easier said than done. On deeper examination, *all* resonant topologies are not similar or identical in all respects. Not all “naturally” offer soft-switching for example. In the well-known words of Bob Mammano (quoted from his presentation titled “Resonant Mode Converter Topologies”, Topic 6, in the 1988/89 series of Unitrode Power Supply Design Seminars): *“While basically simple, this principle can be applied in a wide variety of ways, creating a bewildering array of possible circuits and operating modes”.* In fact not all resonant topologies and modes are necessarily even conducive to the task of reducing switching losses, reducing stresses, improving efficiency, or even reducing EMI. In fact, some do not even lend themselves to a proper *control strategy* as we will see. This last aspect is extremely important but often overlooked in a virtual lab. The entire topic of resonant power conversion turns out to be an extremely complex one: to a) understand, and b) implement effectively.

But do we at least fully understand *conventional* power conversion well-enough? We may realize we need to do better in that regard, because some *critical hints for the successful analysis and implementation of resonant topologies are contained in conventional power conversion*. We may have missed the signs. For example, a potentially puzzling question in a conventional converter is: Why do we ***not*** have any overlap of voltage and current *across the catch diode*? As indicated in Figure 1, there is in fact no significant V-I overlap across the diode, which is the reason we typically assume almost zero switching losses for the diode in an efficiency calculation. This does remain a valid assumption to make, even when we move to synchronous topologies ---- in which we replace (or supplement) the catch diode with *another FET*. So now, consider this puzzle: We now *have a totem pole of near identical FETs (in a synchronous Buck for example), yet we somehow still disregard the switching losses in the lower FET, but not in the upper FET*. How come? In what way is the *location* of the lower FET so different from that of the upper FET? Why does nature seem to favor the top location from the bottom one?

Even more surprisingly, if we delve deeper we will learn that in one particular operating mode, for just *part of the cycle*, even the synchronous Buck exhibits almost no crossover loss in the top (control) FET, and actually shifts those losses to the lower (synchronous) FET! How did this role-reversal happen? *Once we understand all this, we will understand the intuitive direction we need to take in designing good resonant topologies too.*

Understanding and using LLC Converters to Great Advantage

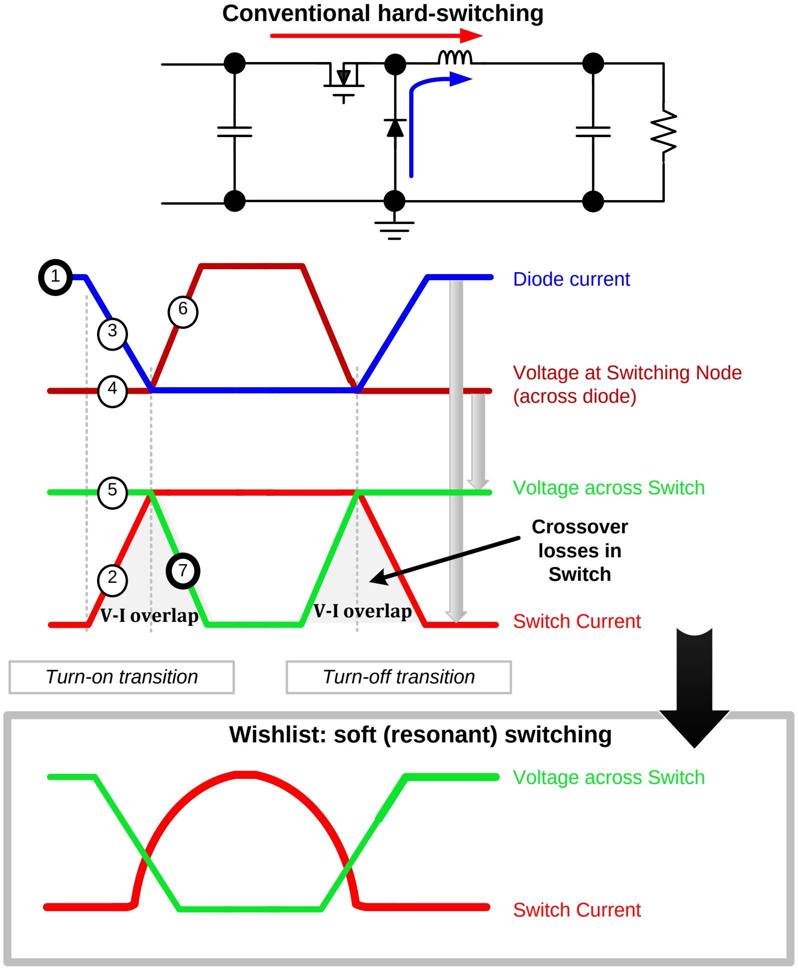

Figure 1: Hard switching compared to resonant (soft) switching

Two Key Concerns (Guiding Criteria)

These form our guiding criteria towards identifying suitable resonant topologies. The most basic questions we need to ask of any proposed circuit are:

Question 1: Does it really reduce switching losses? If so, in which *region of its operation* (or operating mode)? We need to know the line and load boundaries of any such soft-switching region, and thus try to ensure that at least at max load we can significantly reduce switching losses (and preferably continue to do so at light loads too). Finally, we hope we can do this, without asking for, or somehow creating, too wide a variation in the switching frequency. Because, though we may have soft- switching, if we have a very wide variation in switching frequency, that too cannot be good from the viewpoint of either the economics or the size of the EMI filtering stage. We do not gain by simply shifting the cost and dissipation, from one circuit block to another: the EMI stage in this case.

Question 2: The simplest, seemingly most obvious, question can in fact become the hardest to answer: how do we "regulate" the output voltage (of a resonant converter)? Do we do pulse width modulation (PWM) or do something else? And is there a simple way to implement an unambiguous, almost "knee-jerk" (rapid) response to disturbances, *analogous to what we do in conventional power conversion?* We remember that in conventional converters, we pretty much blindly *increase* the duty cycle in response to a *fall* in the output below a set/reference level, for *any reason whatsoever*. Similarly, we decrease the duty cycle if the output rises. That is the heart of basic PWM implementation. We also know that, luckily, for all conventional topologies, a *falling input rail*, or an *increasing load*, *both* tend to cause the output to *fall*, both therefore demanding an *increase* in duty cycle (pulse width) to correct. Though that rapid response does in effect, lead to the complicated area of loop stability, the good news is that we do have a *simple, immediate, and unambiguous response to line/load disturbances*, one which is guaranteed to always be in the *right "direction" for achieving correction*. We do not for example have a complicated control scenario which asks of us something like this: "increase the duty cycle if the input rail falls to 80% of its set value, but thereafter *decrease* the duty cycle (to correctly regulate the output)". This algorithm would be impossible to implement practically. But unfortunately, *in resonant topologies, we can get exactly this odd situation*. If we do not define a very clear region of operation (the allowable/expected line and load variation, related to component selection), we will end up in an area of operation where we are expecting the output to get corrected by our response, but in fact the output veers away in the opposite direction, hopefully collapsing, not overshooting. This becomes another additional, and very tricky, part of the resonant topology puzzle that we seek to uncover in these pages.

Above we have listed the two key concerns or guiding criteria which will be addressed as we go along. But before we do that, as mentioned, first we need to understand a little more about *conventional converters*, and understand in particular, why the switching losses are "lop-sided" --- i.e. losses in one FET, not in the other, in a synchronous Buck converter for example. That will take us to the next step of development of resonant topologies.

Switching Losses in a Synchronous Buck and Lessons Learned

In Figure 1, we first consider the *turn-on* transition (on the left side). Prior to this, the diode is observed carrying the full inductor current (circled "1"). It is obviously forward-biased. But then the switch starts to turn ON, trying to take away the inductor current (circled "2"). Correspondingly, the diode current starts to fall (circled "3"). However, the important point is that while the switch current is still changing, and has not yet taken over the full inductor current, the diode needs to continue to conduct whatever remaining current, i.e. the difference between the inductor current and the switch current, at any given moment. However, to conduct even *some* current, the diode must remain *fully* forward-biased. So, there is no change in the voltage across the diode (circled "4"), or therefore, across the FET (circled "5") as the current through the FET swings from zero to max, and in the diode from max to zero (crossover). Finally, only when the *entire* inductor current has shifted over to the switch, does the diode "let go" of the voltage. The switching node is released, and it flies up very close (a little higher) to the input voltage rail (circled "6") --- and so now, the voltage across the switch is allowed to fall (circled "7") by Kirchhoff's voltage law.

We therefore see that *at turn-on, the voltage across the switch does not change till the current waveform has* ***completed*** *its transition*. We thus get a significant V-I overlap in the switch (FET). That is "hard-switching" by definition.

If we do a similar analysis for the turn-off transition (right side of Figure 1), we will see that for the switch current to start decreasing by even a small amount, the diode must *first* be "positioned" (in terms of voltage) to take up *any* current coming its way (relinquished by the switch). So the voltage at the switching node must *first* fall close to zero (a little below ground), so as to forward-bias the diode. That also means the voltage across the switch must first transition *fully*, *before* the switch current is even allowed to decrease.
We therefore see that at turn-off, the current through the switch does not change till the voltage waveform has **completed** its transition. We thus get a significant V-I overlap in the switch (FET).

But in neither transition, turn-on or turn-off, is there any significant V-I overlap across the diode. We therefore typically assume that there are negligible switching losses in the diode in a conventional topology --- it is the switch (FET) that gets hard-switched during *both* transitions.
Note: In a Boost topology, despite no measurable V-I overlap term present across the diode, the diode can surprisingly cause significant additional switching losses to appear in the FET rather than in itself. For example, we can get losses *in the FET due to* shoot-through (poor reverse recovery characteristics of the catch diode, or the relatively "poor" body diode of a synchronous boost converter's synchronous FET). In other words, the diode may technically never "see" any significant V-I crossover loss across itself, but can certainly cause, or instigate, the switch (FET) to experience significant, perhaps even externally invisible, losses. This particular situation is often tackled by trying to achieve zero-current switching (ZCS), rather than the more common ZVS (zero-voltage switching) which are focusing on in these pages. One simple way out, especially for low-power applications, is to run the converter in DCM, because if the boost diode is no longer carrying current, it has no reverse recovery issues either (it has already recovered when the switch tries to turn ON).

However, now let us look at the synchronous Buck in Figure 2 more closely, at what really happens when the inductor current is *momentarily negative* (below zero, i.e. flowing from top to bottom, through the lower FET). Suppose during this negative current phase, the lower FET turns OFF, while the top FET is forced ON. It turns out we have to pay close attention to what happens *during the deadtime*, i.e. during the small interval between one FET turning OFF and the other turning ON. Most engineers are aware that the primary purpose of deadtime is to avoid any brief possibility of both top and bottom FETS being ON at the very same moment, which would cause efficiency loss due to cross-conduction/shoot-through, and possible destruction of the FETs too. But there is another, subtle advantage that deadtime provides, as explained in Figure 2.

From Figure 2 we can conclude that if the current is momentarily *negative*, and we turn OFF the lower FET at that moment, the inductor will in effect resonate with the Drain-Source parasitic capacitances of the FETs, causing the voltage at the switching node to rise, eventually causing a flow of current back into the input rail via the body-diode of the top FET. Keep in mind that to cause the body diode to conduct, it must be forward-biased. So the voltage on the switching node must swing. Now when the top FET turns ON, it does so with near-zero voltage across it (a forward-biased body diode drop across it). That is “ZVS” by definition. We see the Top-FET being “soft-switched” for a change. Though this effect is occurring only during part of the switching cycle, it is exactly the way it can be made to happen in both (or more) FETs in resonant power conversion too.

Note: Hypothetically, if the average of the inductor current of the Buck falls *completely* below zero (no part of it positive), we would be now constantly drawing current from the “output” (now really an input), and delivering it to the “input” (now really an output). We would have in effect a full-fledged Boost topology, not a Buck anymore. And we also know that in a Boost, it is the *lower* FET that is the “control FET” and the upper FET is the synchronous FET. We therefore intuitively expect to see only the lower (control) FET to have switching losses in a Boost. And, quite expectedly, all the crossover losses of this reversed Buck (which actually makes it a Boost), will now shift to the lower FET over the entire switching cycle. We will have no crossover loss anymore in the top FET. So the role-reversal now makes sense. When only part of the inductor current is negative (in a synchronous Buck), a switch transition during that negative-current time will produce switching losses in the synchronous (bottom) FET, not in the upper (control) FET. When a transition occurs during the positive-current part of the cycle, the switch loss is in the control (top) FET as usual.

We have learned that if we can achieve ZVS for whichever FET has current flowing through its body diode (or an anti-parallel diode placed externally across it) during the preceding dead time (just before transition). So, the *direction* of inductor current at the moment of transition is the key that identifies (or distinguishes) which FET position receives soft-switching and which one gets hard-switched. Keep in mind that what constitutes a “positive” or a “negative” current depends completely on what we are calling the “input” and what is the “output”, and which direction we consider “normal” energy flow.

Note that in either scenario above, we ***certainly need current passing through an inductance to try and “force matters”***. And of course, we also need to leave a small (but not too small) an intervening deadtime for the induced voltage to be able to act. We also need enough inductance to force a *full* voltage swing. Very similarly, in resonant topologies too, whatever the type of simple or complex L-C network being switched, we learn that ***the tank circuit needs to appear “inductive” to the input source, for being able to achieve ZVS***.

Summarizing, the two basic prerequisites for achieving ZVS are:

We need the current to “slosh” *back and forth* --- something that occurs naturally in *resonant* topologies, but is also enforced in conventional half-bridge, full-bridge wave and push-pull topologies. All these are suitable candidates for ZVS, subject to the following stated condition.

The tank circuit (network) impedance must appear *inductive* to the input voltage source, because that is what is responsible, eventually, for trying to brute-force the current through (using induced voltage), in the process creating zero voltage switching across the FETs --- though as mentioned, we also must have *enough inductance* (vis-à-vis available deadtime and parasitic FET output capacitances), so that the voltage on the swinging node can be forced to swing in a *timely* manner (before deadtime runs out), so to reduce the voltage across the FET *before* the FET is turned ON.

Note: Even in ZVS-eligible conventional topologies such as the full-bridge, we can *enforce the second requirement* above, and produce "quasi-resonant (QR) ZVS". These topologies are all based on the Forward converter topology using a transformer. We need to add a small Primary-side inductance (often just the transformer leakage) for this purpose.

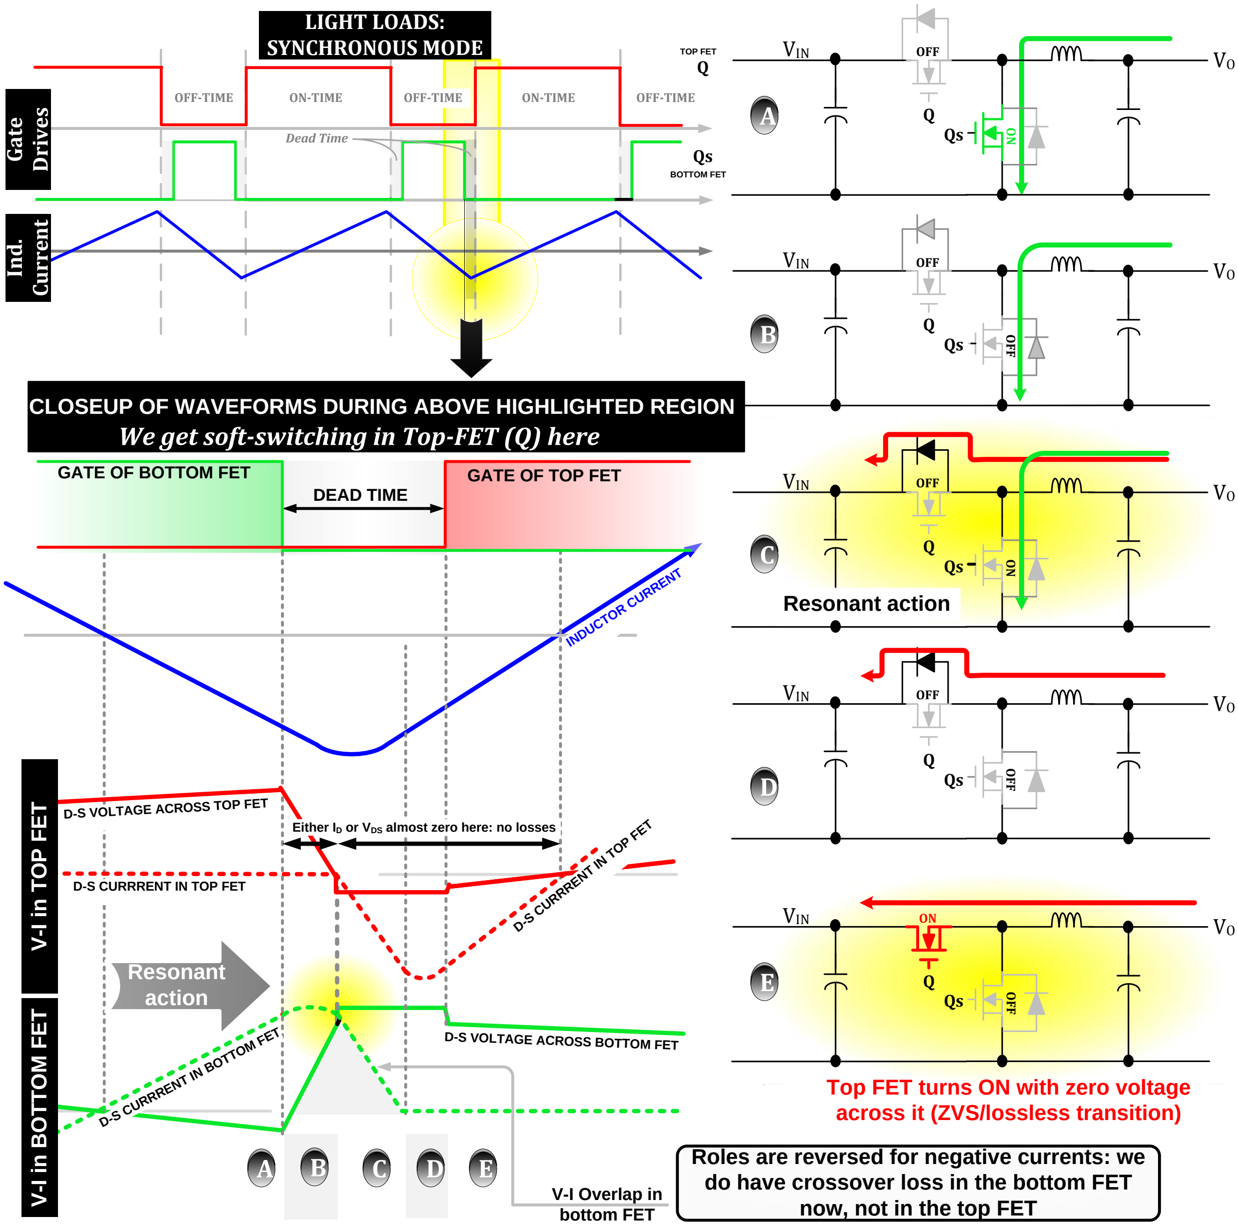

Figure 2: We do get ZVS naturally, in some cases/modes of the synchronous Buck too

Building Basic Resonant Circuits from Components: the "PRC"

In resonant topologies, just as in conventional switching power conversion, we try to pick *reactive* components (inductors and capacitors). Because we know that ideally, they store energy, but cannot dissipate any. We always require *resistance*, either in the form of intervening parasitics, or in the form of the load itself, to dissipate any energy, either usefully or wastefully. But in addition, the L and C can pass energy back and forth *between each other* on account of their *complementary phase angles*. We can try to use that useful property to transport energy on their shoulders, from the input (source) to the output (load), somehow creating *DC regulation too along the way if possible,* which is a fundamental requirement of any type of practical power converter.

That is the general direction for implementing any "lossless" (high-efficiency) power conversion methodology. The key difference between resonant topologies and conventional switching topologies is primarily related to the actual "values" of the L and C components used. In conventional power conversion, we use relatively "large" capacitors at the input and output, so the natural LC frequency happens to be very large relative to the switching frequency. We just do not "see" resonant effects anymore, because *we do not wait that long* before we switch. But if we reduce the L and C values significantly, we will start to see resonant effects even in conventional power conversion. Keep in mind though, that just because they are "resonant", doesn't make them *acceptable* or useful. Their eventual usefulness is judged mainly on the basis of the two "guiding criteria" described previously.

Note: Modifying conventional topologies by reducing their L and C values is usually that is not the best way to go in creating resonance, except for example in the "ZVS phase-modulated full-bridge", which is best described as a "crossover converter" (in the sense of a crossover-vehicle category), literally bridging conventional power conversion with resonant effects (but resonance occurring only *during* the transition deadtimes). This maintains the desirable characteristics of constant clock frequency of conventional PWM topologies, but uses resonance during deadtime for inducing soft-switching (in all the FETs).

To create proper resonant circuits, let us start with a simple inductance and a simple capacitance. But they are *not connected to each other yet*. To make this "real", let us pick some numerical values. We choose L = 100mH and C= 10µF. These may seem rather big, but as a result, our switching frequency is also going to be low, and that is quite helpful initially at least, for ease-of-discussion, and for cleaner and faster circuit simulations etc. So, everything is initially being scaled to a lower switching frequency just for convenience.

Let us connect each separated L and C component to *identical* AC sources, of 30V (amplitude, or half peak-to-peak value) with a frequency of 300Hz, and just see the currents through each. We run this through a Simplis simulator. The results are presented in Figure 3.

As expected the input and output voltages on either component are the same, at 30V. The currents are different. We expect the corresponding current amplitudes to be:

$Z_C=1/(2\times\pi\times f\times C) = 1/(2\times\pi\times300\times10^{-5}) = 53.05\Omega$. So $I_C = V/Z_C = 30/53.05 =$ **0.565 A**.
$Z_L= (2\times\pi\times f\times L) = (2\times\pi\times300\times0.1) = 188.5\ \Omega$. So $I_L = V/Z_L = 30/188.5 =$ **0.159 A**.

This is exactly what we got through the simulations shown in Figure 3. But via that figure, we see another possibility, one that we can hopefully exploit: the peak of the inductor current comes a little *later* (exactly $1/4^{th}$ of a cycle) after the peak of the inductor voltage (which is the same as input voltage in this case). That is why we usually say that the *current lags the voltage in an inductor by 90°*. We can confirm from conventional switching power conversion too, that when we apply a voltage across an inductor, the current ramps up slowly. So current does lag the voltage in that sense, though unfortunately, we can't really visualize or define what the "lag" is for non-sinusoidal waveforms as in conventional power conversion. Now, looking at Figure 3 once again, we see that the cap voltage peak *lags* the cap current peak by exactly the same amount, i.e. 90°. In other words, the ***capacitor and inductor currents are relatively exactly 180° out of phase.*** They are "complementary", because 180° is *just a change of sign* or direction (with respect to each other). Note that we have implicitly chosen/assumed the convention that current *into* the component (L or C) is "positive", whereas current coming out of it is "negative". So a 180° relative phase shift simply means that when 159 mA is coming *out* of the inductor, *exactly at that moment*, 565 mA is going *into* the capacitor. We can confirm this from Figure 3.

The preceding logic also leads us to the following thought process: Could we have "X" mA coming out of the inductor and exactly "X" mA going into the capacitor at the same moment? Yes, by varying the frequency we can always do that --- because in one case (inductor) the impedance increases with frequency whereas in the other case (capacitor) it decreases with frequency. So certainly, we can get the two values to converge at some "intersection" point, at which point their impedances would be equal *in magnitude* (opposite in sign, though). See Figure 4. That intersection then forms a "natural (or resonant) frequency" for the chosen L and C. The only question is: at what frequency? That is just the point at which the inductor's impedance $2\pi f\times L$, equals the capacitor's impedance $1/(2\pi f \times C)$. Equating the two, we solve to get the well-known equation for "resonant frequency": $f_{RES} = 1/\{2\pi\times\sqrt{(LC)}\}$. We will describe resonance more clearly now, based on the above observations. We could connect the two components (L and C) in parallel across each other, inject some energy into the "tank", say by tuning the applied AC source to the natural frequency of the two. Thereafter, *theoretically speaking*, once injected into this tank, energy could just slosh back and forth forever between the L and the C. It would be self-sustaining. See Figure 4. See also Figure 5 for a more detailed explanation of the phenomenon.

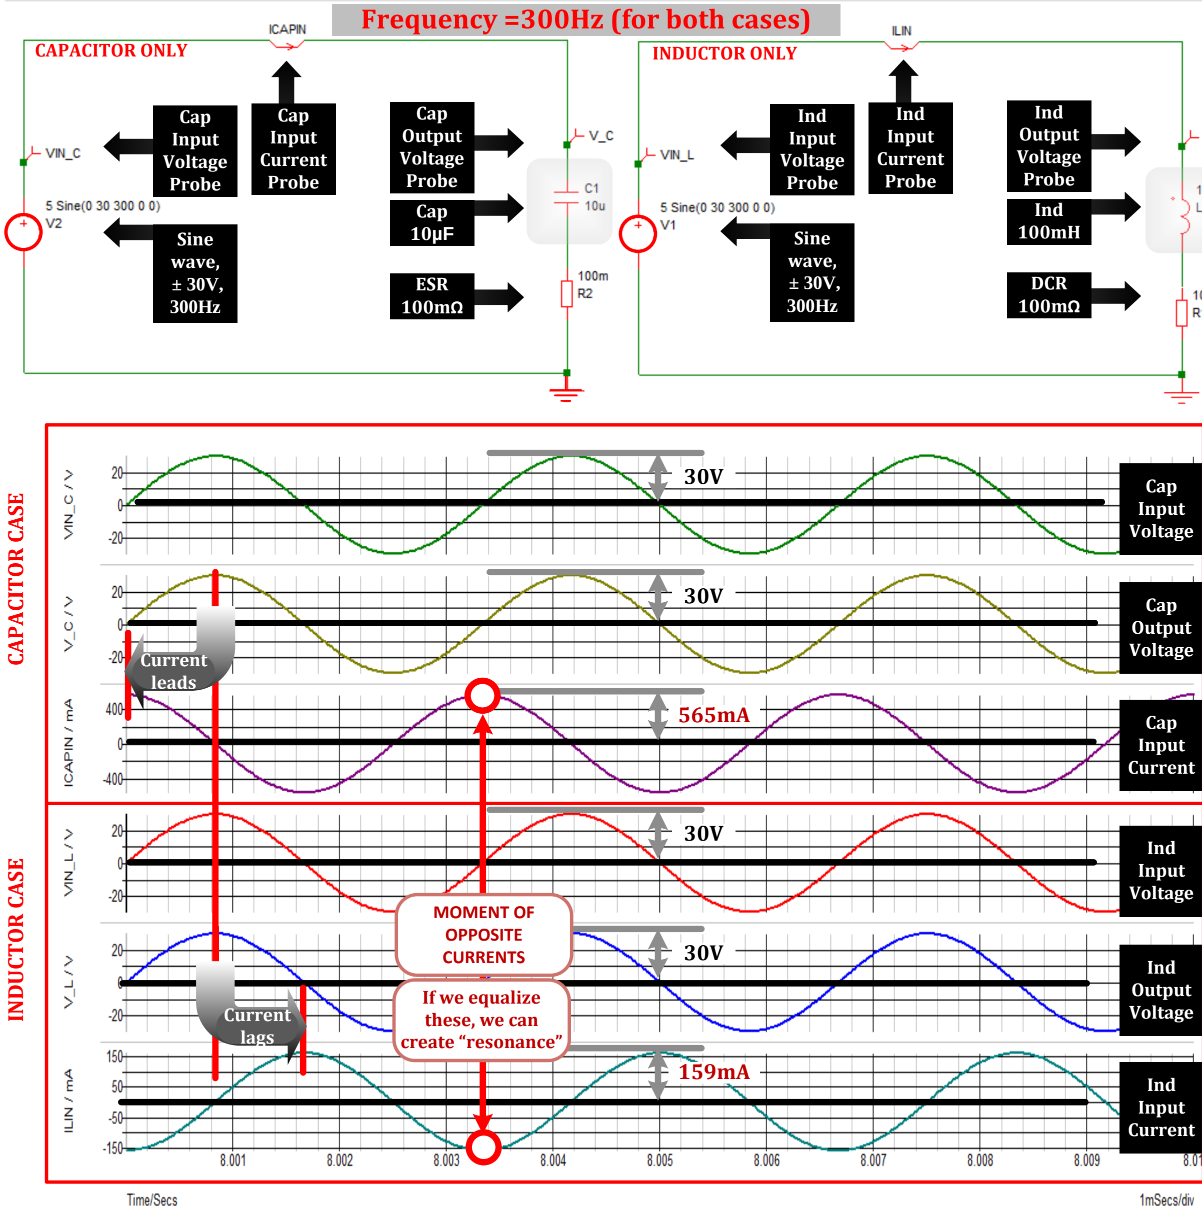

Figure 3: Response of a 100mH inductor and a 10µF capacitor to identical 30V/300Hz AC sources

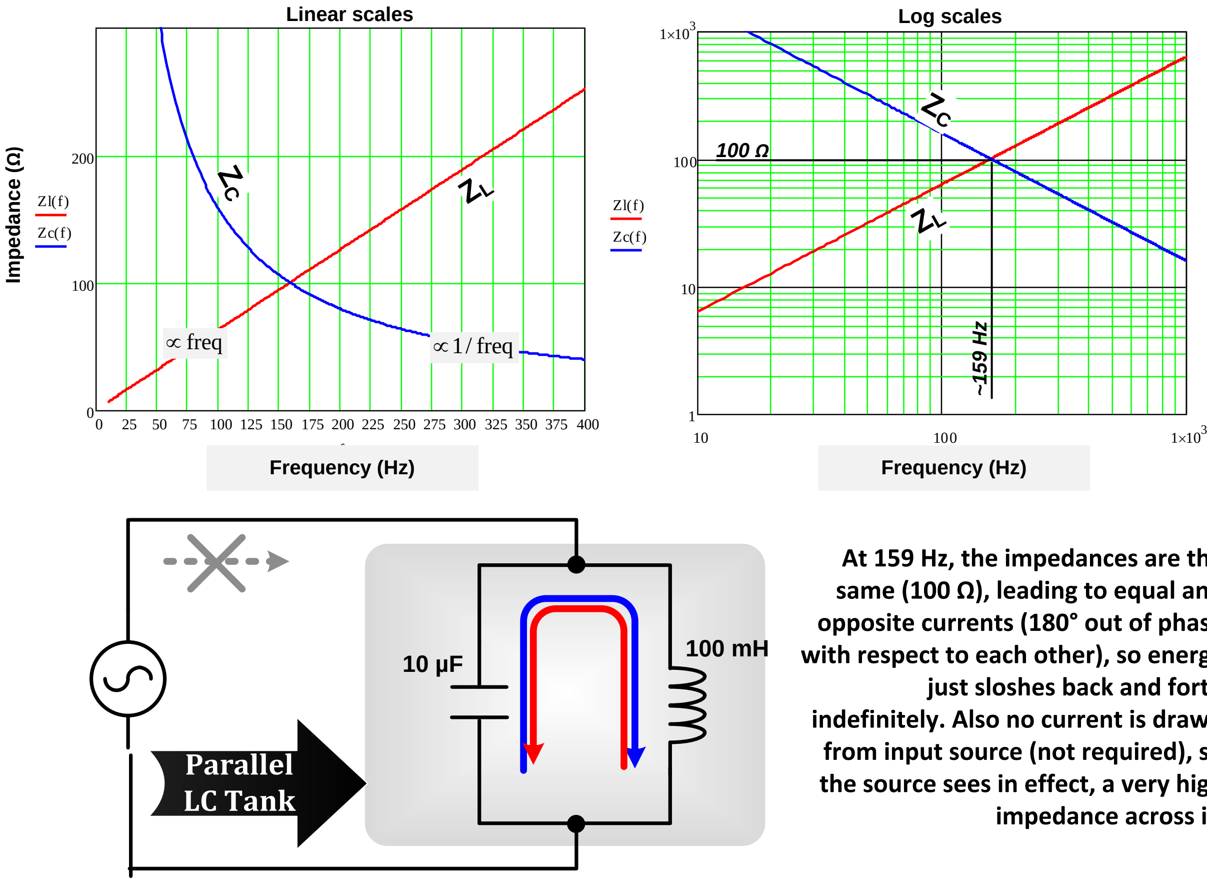

At 159 Hz, the impedances are the same (100 Ω), leading to equal and opposite currents (180° out of phase with respect to each other), so energy just sloshes back and forth indefinitely. Also no current is drawn from input source (not required), so the source sees in effect, a very high impedance across it.

Impedance of the LC at resonance is not 100 || 100 = 50Ω. It is infinite!

Figure 4: Explaining the energy storage capabilities of the parallel resonance tank circuit when driven at its resonant (natural) frequency

Understanding and using LLC Converters to Great Advantage

PARALLEL LC AT RESONANCE

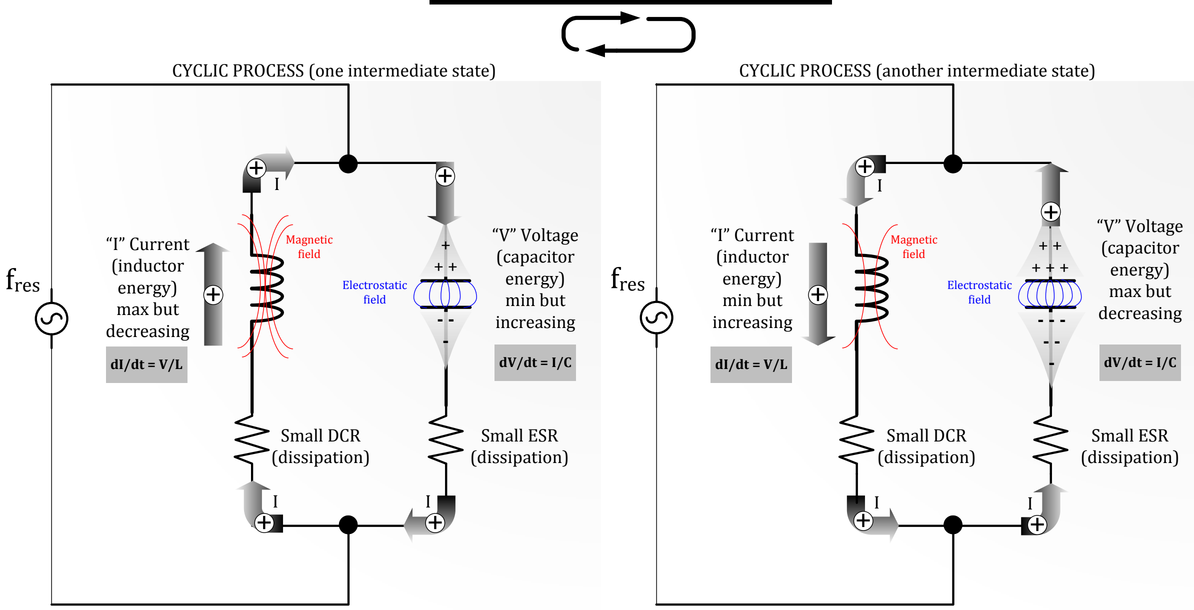

By physically paralleling the two components, L and C both are forced to have the same magnitude of voltage across them at any instant. At resonance, both have the same impedance too. So their **currents**, though equal, have a 180° *relative* phase shift --- i.e., the instantaneous current *out* of either one, almost exactly equals the current *into* the other (difference is due to non-idealities). So, during "resonance", energy (current) sloshes back and forth between inductor and capacitor as shown above. When current is peaking (i.e., inductor energy is max), its rate of rise (i.e. voltage, and therefore cap stored energy) is very small. When the voltage is peaking (i.e., max cap energy), its rate of rise (i.e., current, and therefore inductor stored energy) is very small. That leads to the back-and-forth energy sloshing. **Huge out-of-phase currents can be generated in the L and the C separately, but they almost cancel out, so the voltage source does not "see" those huge currents.**

This sloshing can go on forever, were it not for real-world parasitic resistances such as ESR and DCR. So, a small amount of energy is constantly dissipated during the sloshing, and the same gets pulled in from the voltage source to compensate for this loss and to maintain a steady state (provided that a voltage source is present, because otherwise the LC oscillations will decay exponentially). But being a parallel circuit, the voltage across the L and the C equals that of the voltage source, with a relatively small current (from the voltage source) passing through the LC combination . Since this small input current is only drawn for compensating a small purely resistive loss, the input current is in phase with the input voltage at resonance. To the input voltage source, the parallel-LCR network appears as a simple resistance (of a typically ***large value***, producing ***very small input currents***), at the resonant frequency.

Since the current drawn from the voltage source is very small, we can say that the **parallel LC circuit offers a very high net impedance at resonance**.

A small (large) resistance in series with the L for example, can be treated as a high (low) resistance across an ideal L. Though the transformation (equivalent parallel resistance) is frequency dependent. Similarly, a small (large) resistance in series with the C for example, can also be treated as a high (low) resistance across an ideal C. Though the transformation (equivalent parallel resistance) is frequency dependent again.

Figure 5: Explaining the energy storage of the parallel LC in more detail

Note that once the LC "tank" has been excited and "set in motion", we could even remove the AC source completely, and the oscillations would continue indefinitely (ideally). But note that if we had continued to keep the AC source connected to this tank circuit, it would make no difference at all really (unless of course the AC source tried to drive the tank circuit at some frequency *other* than its natural frequency). At the resonant (natural) frequency, provided the LC circuit has reached the voltage level of the AC source, the AC source does *not* need to replenish the energy in the tank (assuming no parasitic series resistances present). In other words, the AC source will thereafter provide *no* current at all, even if connected --- simply because it does *not need to*. But we also know from our usual electrical definitions, that if the current drawn from the input source is zero, then the *impedance (of the LC tank connected across the source), i.e., as seen by the source, is infinite (just like an open circuit, though only true at that specific frequency)*.

Above we have created our first resonant circuit: a pure "LC" parallel circuit (with no parasitic resistances). It is our first building block. We have intuitively also just learned that this pure parallel LC tank circuit has an infinite impedance at its resonant frequency "f_{RES}".

As mentioned, we have the result: $f_{RES} = 1/\{2\pi\times\sqrt{(LC)}\}$.

It is interesting to point out that at resonance, the overall impedance of the tank circuit is not half the impedance of each equal limb as we expect in the case of paralleled resistors. That is because of the complementary phase angles between voltage and current in the paralleled components, that the net impedance becomes infinite in magnitude, at the resonant frequency.

In reality, any small resistances in series with the C and the L (such as ESR and DCR), will cause the oscillations to decay exponentially. And in that case, to "replenish" the tank, the AC source will need to provide a small amount of current, *even at resonance*. Which implies that though the impedance is still very high, it is not "infinite" anymore, not even at the resonant frequency.

We have simulated an (almost unloaded) parallel LC circuit --- with a 30V AC source set exactly to the natural (computed) frequency of 159.155 Hz for the selected components (10µF and 100mH). The results are presented in Figure 6. Compare that with Figure 3 where we still had "unmarried" and separate components. If we had not loaded the circuit a bit, the Simplis simulator would have complained due to infinite numbers.

The impedance of the parallel LC falls off on either side of the resonant frequency, irrespective of parasitics being present or not. At very low frequencies, we can assume the inductor is just a short (piece of wire), and so the AC source too will see a short (zero impedance), whereas at very high frequencies, the capacitor becomes a short (high frequency bypass), so again, the AC source will see almost zero impedance.

We sum this up by saying that to the ***left of the resonant frequency of a parallel LC circuit, the LC network appears "inductive"*** (current lagging the voltage, but not necessarily by a full 90°). Whereas to the right of the resonant peak, the network appears "capacitive" (voltage lagging the current, but not necessarily by a full 90°). Since we have learned in the previous sections that a resonant network or circuit should appear "inductive" to the source, to be able to create ZVS, we realize that ***the parallel LC circuit is useful to us only provided we operate to the left of its resonant peak***.

Any small resistances in series with the L and C (their typical parasitics like ESR and DCR) can be transformed, or modeled, into an equivalent large resistance placed in parallel to the paralleled (pure) L and C. In any proposed resonant converter based on this type of tank circuit, the load too will be connected in parallel to the paralleled L and C. So the effect of all resistances is to cause a eventual decay of the stored energy if the AC source is suddenly disconnected. If the AC source continues to be connected, it now has to constantly provide current and energy into the system to keep it "topped up" --- in the form of a) useful energy delivered to the load, and b) wasted energy dissipated in the series parasitics (or the equivalent parallel resistance).

Note: The "series to parallel equivalent transformation" indicated above, in effect, makes the parallel resistance a function of the AC frequency. For any given frequency we have to recalculate it.

Finally, keep in mind that though the AC source may be providing only a tiny current at resonance, the current sloshing back and forth in the parallel LC can be *very high*. It is limited only by the impedances of the individual L and C, as we can see from Figure 6. The AC source may never "know" about these high currents since it is only delivering a very small current, but in a practical converter, where we would have transistors in the path of the circulating current, we are in danger of damaging our PRC (parallel resonant converter based on this building block tank circuit). That is one limitation. Another limitation of this very basic PRC is that there is no obvious way to regulate the output! The load is connected in parallel to the input source, so it has exactly the *same* voltage as the incoming rail. What use is that if we can't offer load and line regulation? Yes, we can add other L and C's to try to get this to work, but at this point we prefer to move on to another, more promising resonant converter variation, which was in fact quite popular in older resonant converters: the SRC (series resonant converter). It is based on the series resonant LC cell, discussed next.

Note: We can prove that the real-world actually depends on their being "parasitics" present almost everywhere. Very strange things would happen in our "real-world" if there were no *resistive* parasitics present in particular. It turns out, it is also a very good idea to include small resistive parasitics during simulation, just as we have done in Figure 6. To avoid mysterious simulation stalls, we should also try putting in initial conditions for the L and C, for the voltage across the cap or the initial current through the inductor.

PARALLEL LC AT RESONANT FREQUENCY (159.2Hz)

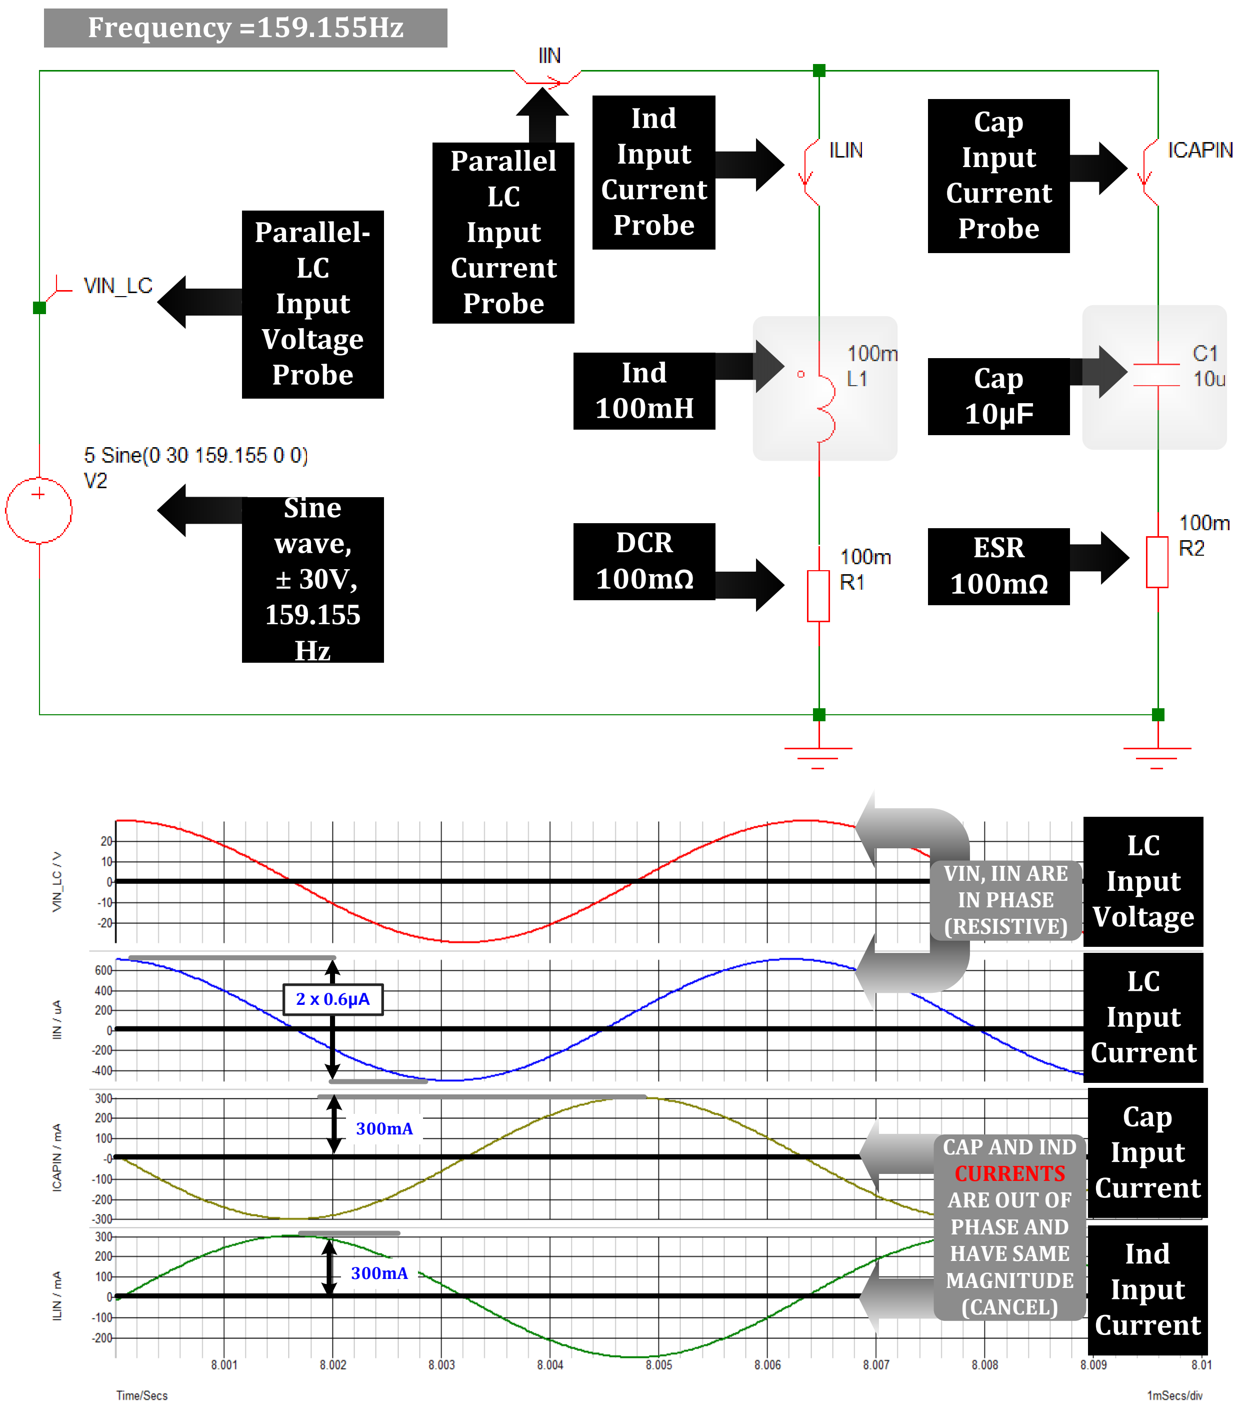

Figure 6: The (almost unloaded) parallel LC at resonance, showing the high circulating currents in L and C

PART 2: Building Resonant Tank Circuits for Future Converters

After having completed the initial transition from "PWM-thinking" to a better understanding of resonant behavior, we now start putting the LC components together to create some "trial balloons" first. Eventually we will then understand the best direction to take for practical converters.

Series Resonant Tank Circuit

This formed the basis of several commercial resonant converters. Let us get familiar with its pros and cons.
In Figure 7, we have simulated the series resonant tank, driven at its resonant frequency, almost fully unloaded (i.e. small parasitic resistances too). We see very *high currents* at resonance, clearly *limited only by the parasitics*. Then using V=I×Z, because of the very high currents, we also get *very high voltages* across the L and C. However, these high voltages are out of phase, and almost fully cancel out from the viewpoint of the input source. So we are just left with 30V (AC amplitude) to satisfy Kirchhoff's voltage law. The AC source does not "see" the high voltages, but does see the very high currents because they pass through it.

In a practical converter based on this series LC principle, we would typically connect the load across the L. Now we can use a "parallel to series equivalent (frequency dependent) transformation" and visualize this parallel load resistor as appearing in series with the L and C, just like the resistive parasitics (ESR and DCR) do. Or we could, in fact, really place the load in series with the L and C as shown in Figure 8. Whatever way, the load helps significantly reduce the high voltages and currents of the series LC (in any proposed SRC, i.e., series resonant converter). This is called "damping". It does make the series LC tank a possible practical choice, *provided we do not try to run it with no load across L*--- because damping would be lost in that case, and we could easily damage it due to the high peaking at resonance of the series LC . That is one limitation of the SRC (based on this tank circuit).

Can we at least use the SRC to provide a regulated output rail, something we couldn't do easily with a proposed PRC? As shown in Figure 8, we in fact use the basic voltage divider principle we use in setting the output voltage of a typical conventional PWM regulator, to produce a voltage rail (always) lower than the input. Nom we have to *vary the impedance of the LC* appropriately to produce whatever output we want. In other words, ***in a series resonant circuit (in general many types of resonant circuits), we need to vary the "switching" (driving) frequency to create output regulation, just as in a conventional converter we need to vary the pulse width to create regulation.***

Can we ensure ZVS in a series LC based converter? For that we need the tank circuit to appear *inductive* to the AC source as explained. So, without bothering to plot it out so far here, we can intuitively visualize that at low frequencies, the inductance of the series LC would appear as a piece of wire in series with a cap, so the AC source would only see a capacitance at low frequencies. At very high frequencies, the cap would appear shorted to an AC signal, so the tank would now appear inductive. We conclude that in a practical SRC, ***we would need to operate to the right of the resonant frequency*** for achieving ZVS.

One famous "last-but-not-the-least" type of question remains: if the output falls, do we need to *increase* the frequency or *decrease* it? To figure that out, in Figure 9, we plot out the gain of the series LC (we call gain as "conversion ratio" at various places, but it happens to be just the output divided by the input). We see that in the "ZVS-valid" region to the right of the resonant peak, we need to reduce the frequency as load increases (to get the gain to increase). Conversely, at light loads, we have to significantly increase the frequency to regulate. For the case of R = 100 Ω, we see we need to go from 300 Hz to 170 Hz to regulate the output (keeping the same conversion ratio, since input/line has not changed here). For R = 1000 Ω, we are already in big trouble: theoretically, we have an *almost infinite switching frequency spread* in an SRC, just to regulate to a set level (at ~ zero load). We also know that at very light loads, the voltages and currents can be extremely high in a series LC tank, which not only can damage the switches of an SRC, but result in very poor efficiency at light loads. These are all the major limitations of the series-LC and its practical form, the SRC.

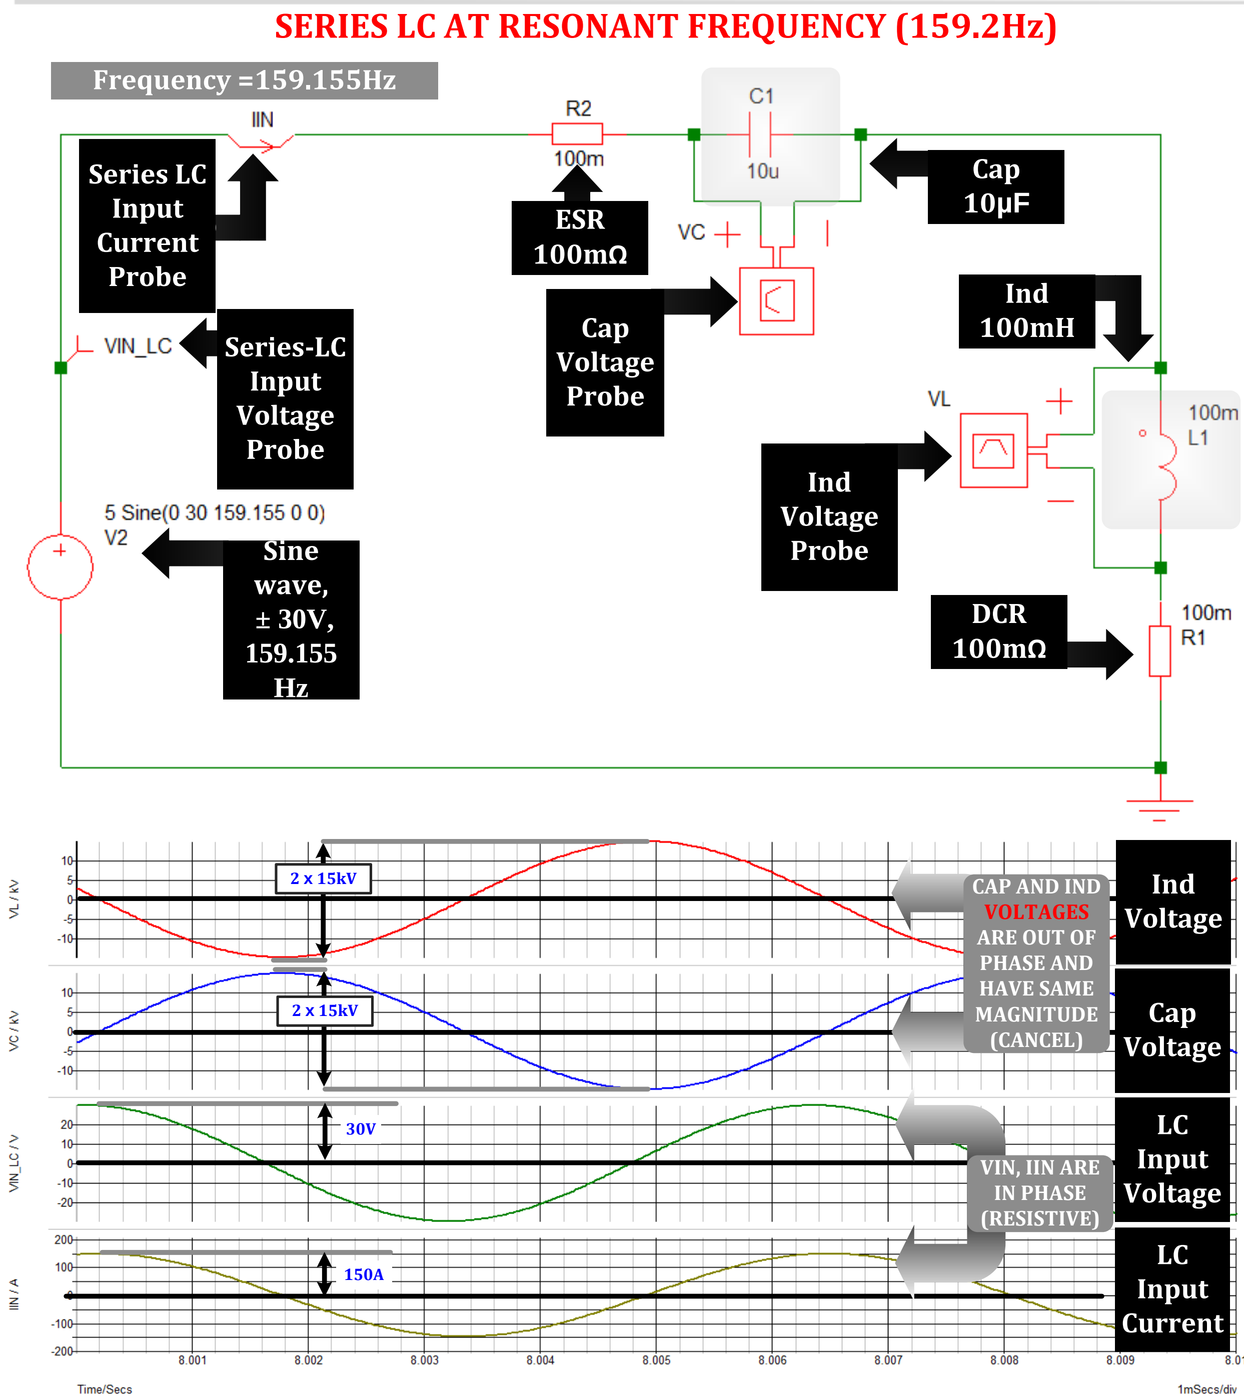

Figure 7: The (almost unloaded) series LC at resonance, showing the high voltages across L and C

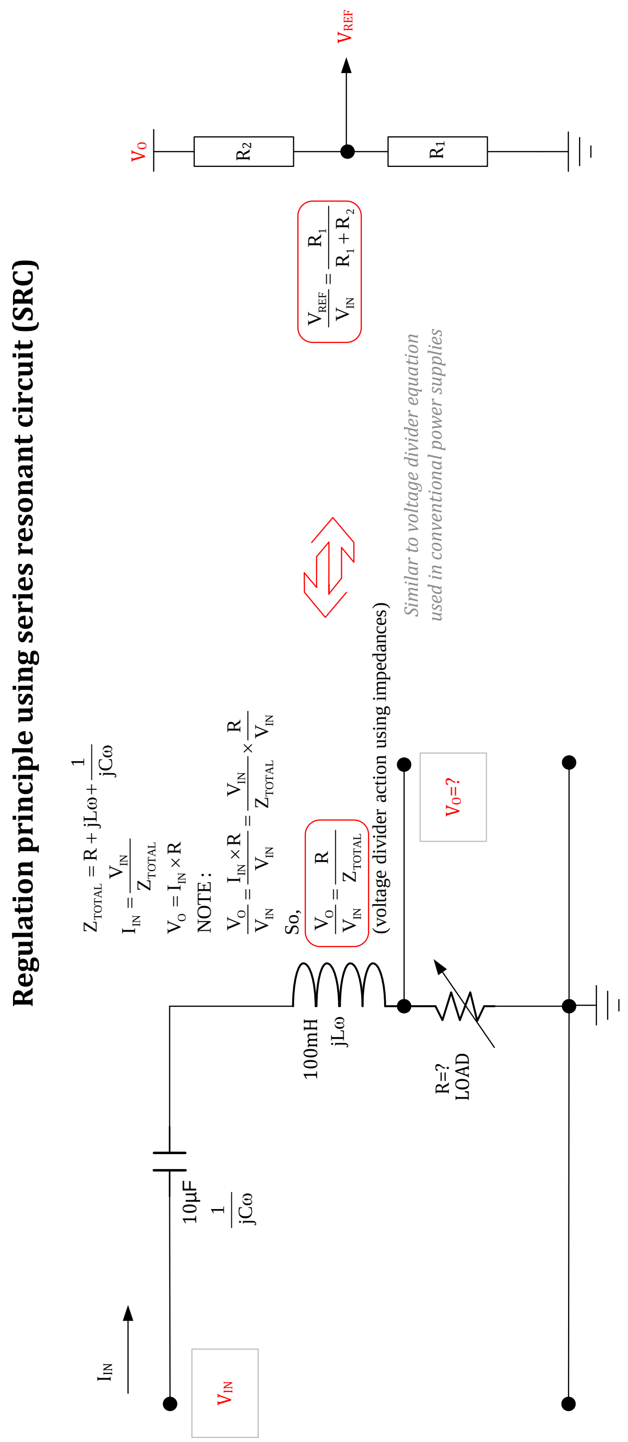

Figure 8: Using the series resonant circuit to create a regulated voltage in principle

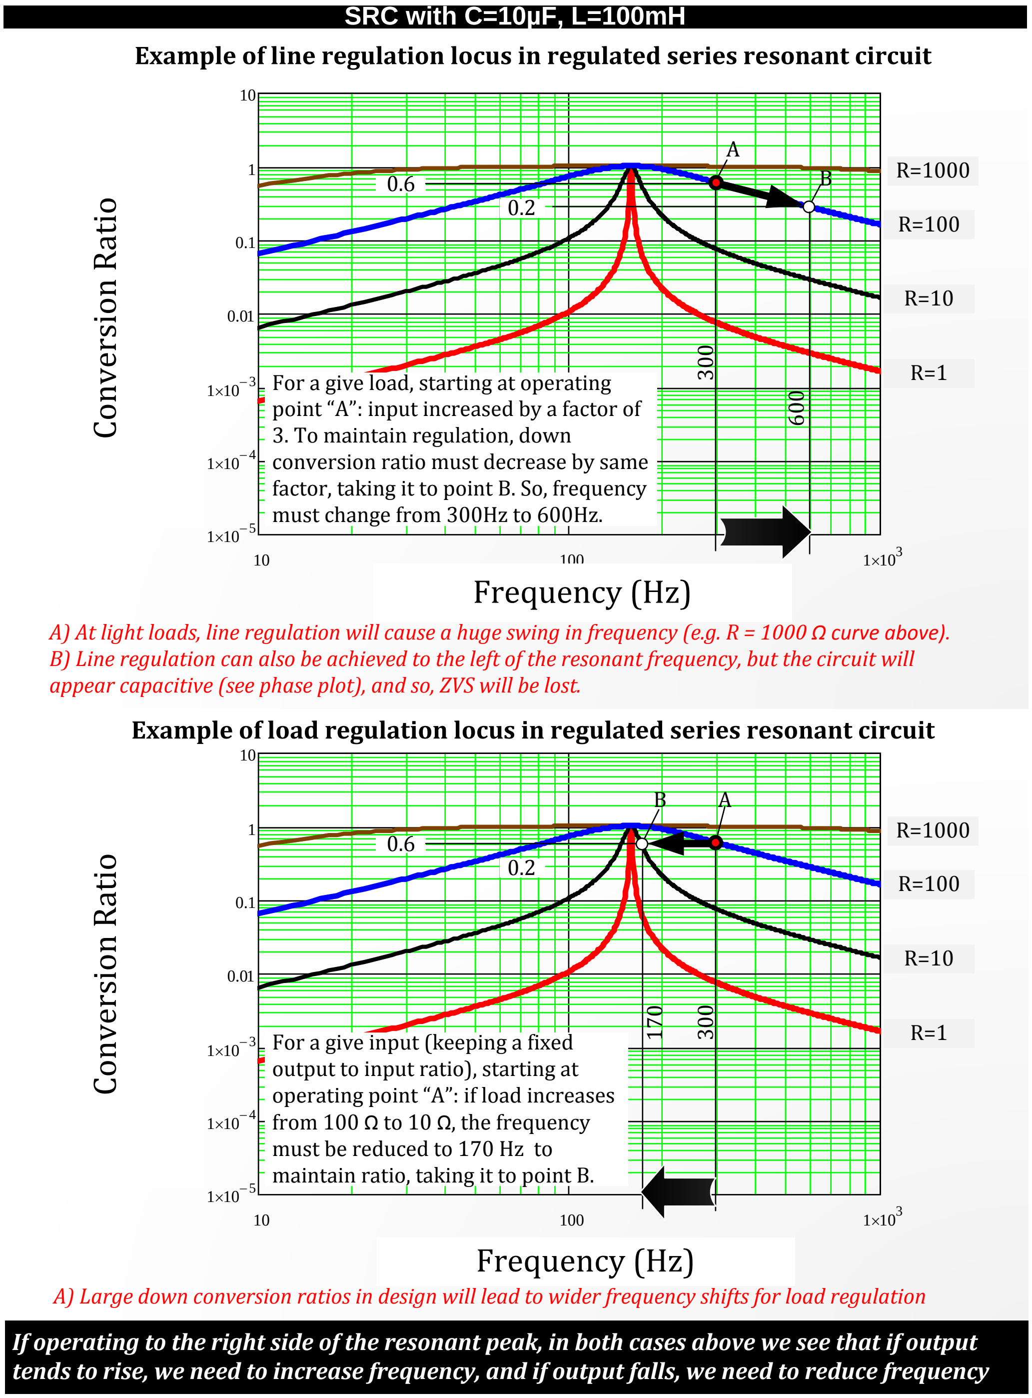

Figure 9: The series resonant LC tank, showing gain (conversion ratios), and how line and load regulation can possibly be carried out in a practical SRC (series resonant converter)

Understanding and using LLC Converters to Great Advantage

Introducing the LLC Tank for creating future LLC converters

Here all we do is to take the series LLC tank and add a *relatively large inductor in parallel to the load*, as shown in Figure 10. So, now we have two inductors and one capacitor. We call this "LLC" for L-L-C (two inductors, one capacitor).

Historically, the LLC was not really an unknown topology. But its full significance was not clearly understood till just a few years ago when engineers started looking rather keenly once again at resonant topologies in an effort to reduce switching losses.

Note: It is paradoxical though that the design of LLC converters is still so poorly understood that even though LLC converters are based on a principle that should help achieve very high efficiencies at very high frequencies, (where the last stumbling block was always the "switching loss" term), yet, most commercial LLC converters are still operating only in the range of 80kHz to 200kHz. Technology will certainly improve with a better understanding. There have been 1 MHz LLC prototypes already reported.

We need to understand very clearly how to properly and optimally design an LLC converter --- *based on physical principles and deeper understanding*, rather than just relying on "trial and error" in a real lab, or on simulations in a virtual lab. It is tricky. As even Bob Mammano admitted: it can be "bewildering". We hope to overcome some, if not all, the mystique behind the LLC converter through these pages.
We expect two resonances, because we have one C, which can "resonate" with not one, but two inductors. Note that two L's cannot "resonate" with each other, because we need complementary phase angles to resonate! We have to see how this additional L modifies the series LC circuit, and if it helps overcome some of the previously discussed limitations of the SRC.

In Figure 10, we show how the voltage divider principle can be made to work here too, this time similar to the case of a *three*-resistor voltage divider which i principle, can be used for setting the output in any conventional PWM converter. Then, in Figure 11, we have used Mathcad to plot the equation introduced in Figure 10. But before we get to fully discussing Figure 11 (the gain), let us first analyze the *phase* relationships accruing from the basic equation in Figure 10 , so we can be sure to identify the region of operation *in which the circuit appears "inductive"* and is therefore conducive to establishing ZVS. We have to admit, we cannot analyze all this very intuitively anymore, and need to plot it out mathematically as shown in Figure 12. Note that as for the series LC circuit, we have chosen L_1 to be 100mH, and C is still 10µF. The newly introduced inductance is L_2, and we have used a seemingly arbitrary value of 900mH (a factor of 9 higher as compared to L_1). Actually, that factor is optimally set. If the ratio is larger, the frequency variations are more, and if it is made too small (similar to L_1), the currents cam be much higher and the efficiency much lower.

In the next section, we start by analyzing Figure 12 and then Figure 11 again, going back and forth to draw pointers towards designing a practical LLC switching converter. Note that for now, *we are only discussing how the **LLC tank circuit behaves**, and that too only under a **sine wave** stimulus (AC source). Later we will build a practical switching converter out of it in several steps.*

Understanding and using LLC Converters to Great Advantage

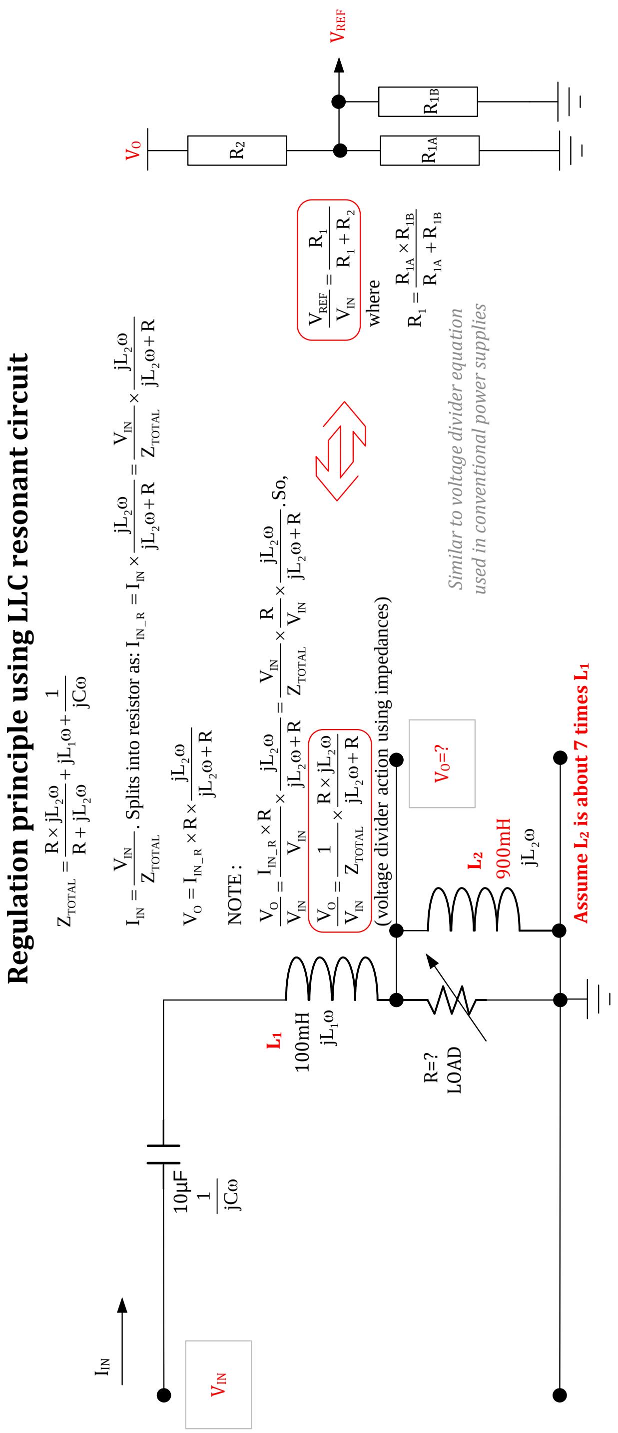

Figure 10: Using the LLC resonant circuit to create a regulated voltage in principle

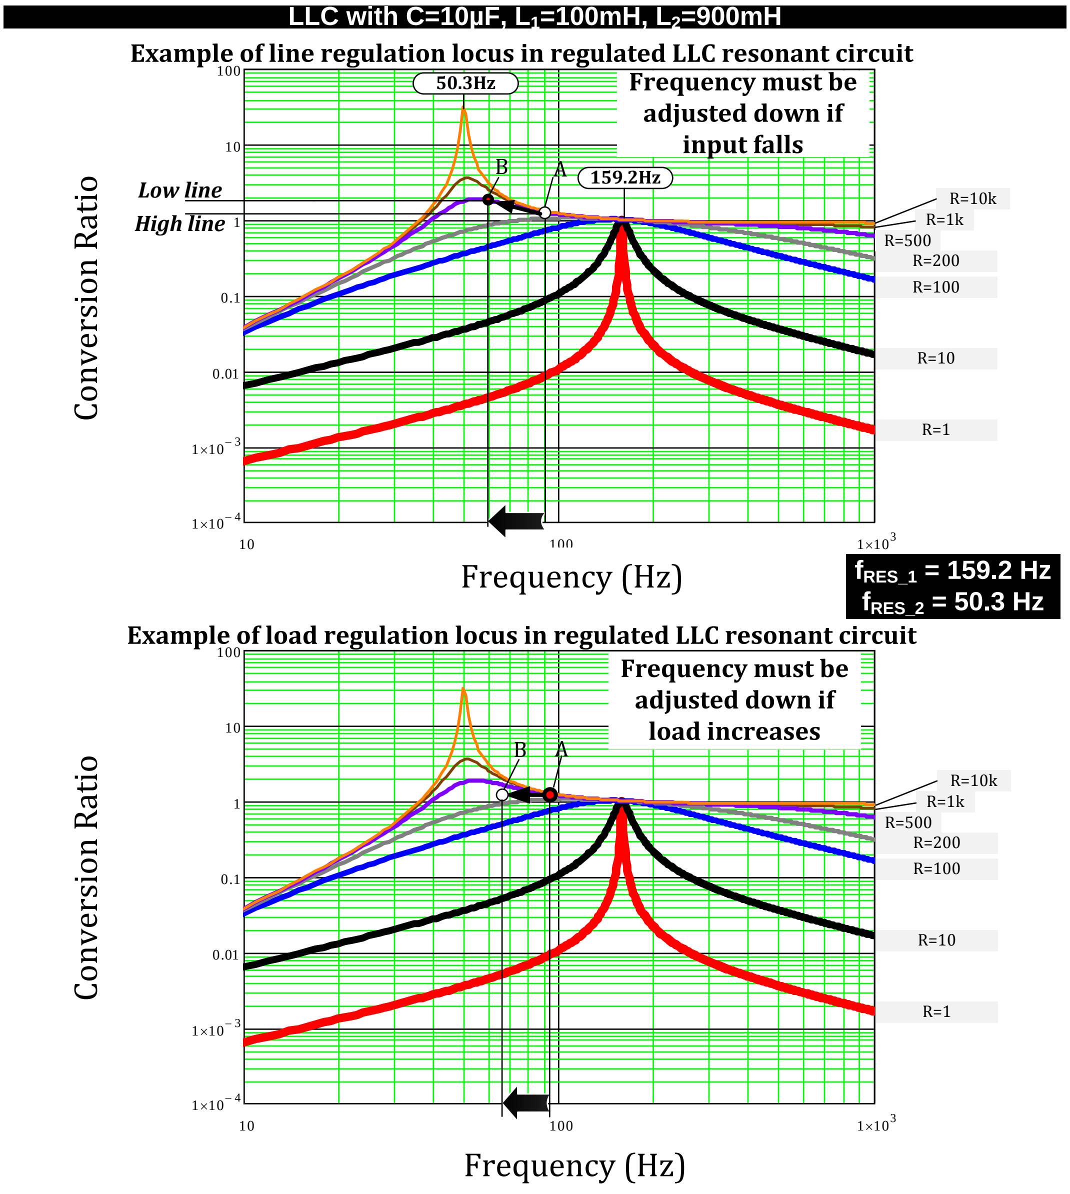

Increasing the load or reducing the input, both tend to cause the output to fall, and thus the regulation loop will need to try and correct this by increasing the conversion ratio (Output/Input). A) If the system is operated ***between the two resonant frequencies*** *at all loads, we can set that the logic that the control loop will* ***always simply decrease the switching frequency in response to a falling output****. B) We can also restrict ourselves completely to the right side of the right-hand resonant peak at all loads. But that has all the disadvantages of conventional series resonant circuits (e.g., huge, uncontrolled swing in frequency at light loads). C) We can also decide to restrict ourselves* ***completely to the left side of the left-hand resonant peak,*** *and set the logic such that the control loop will* ***always simply increase the switching frequency in response to a falling output****. That too will work in terms of regulation too, but it will not be ZVS. D) But we can never set an easy control algorithm if we allow operation across either resonant peak, since conversion ratio slope changes on either side.*

Figure 11: The gain (conversion ratio) of the LLC resonant circuit for different loads

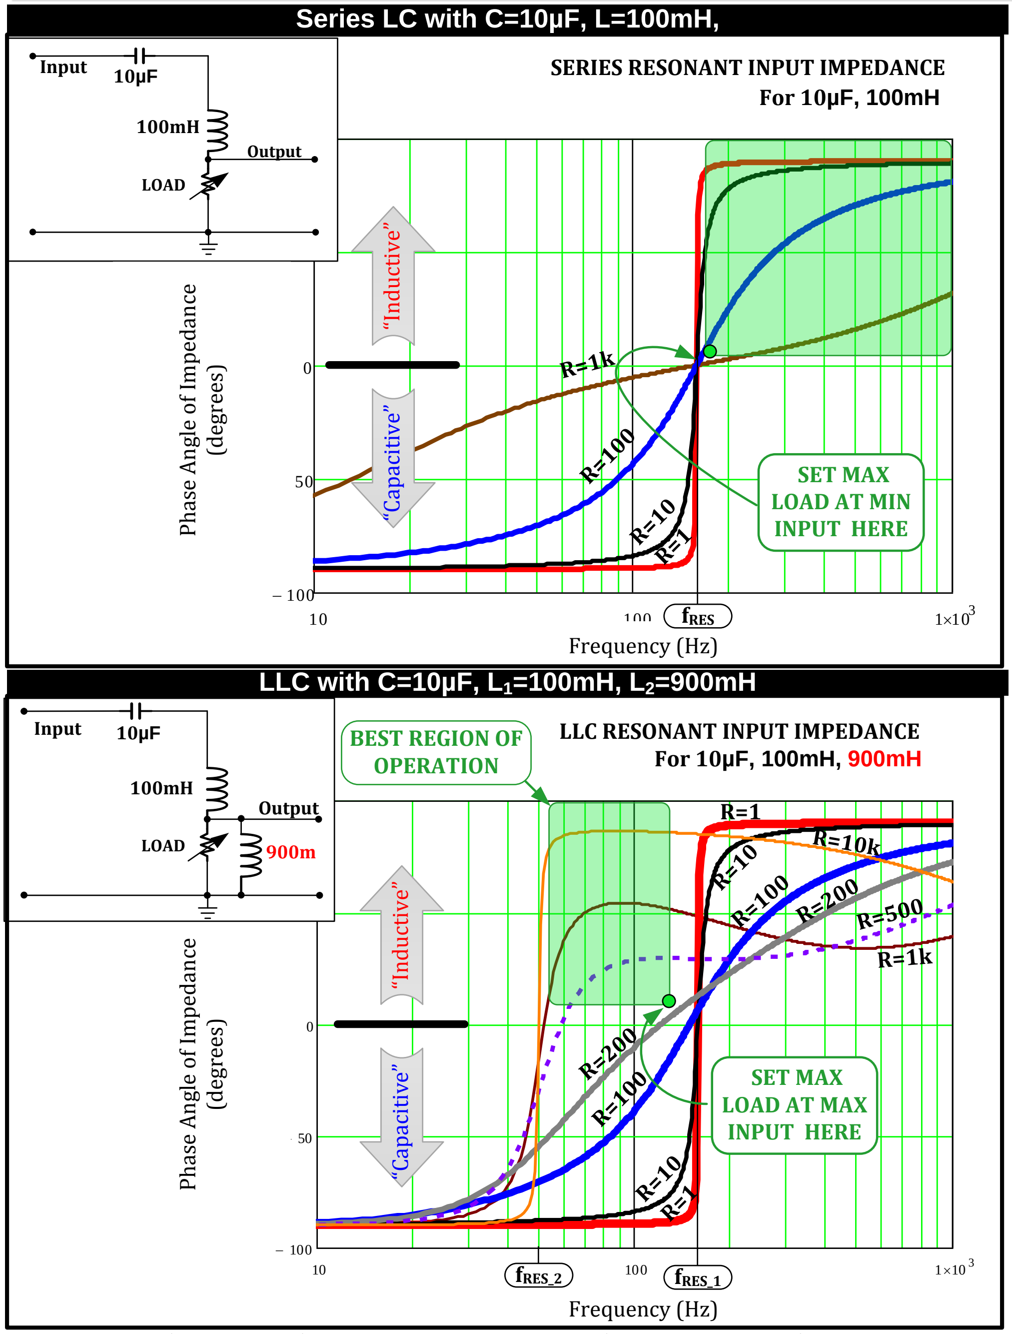

$$f_{RES} = f_{RES\_1} = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{100mH \times 10\mu F}} = 159.2Hz; f_{RES\_2} = \frac{1}{2\pi\sqrt{(L_1+L_2)C}} = \frac{1}{2\pi\sqrt{(100+900)mH \times 10\mu F}} = 50.3Hz$$

Figure 12: The phase of the LLC resonant circuit for different loads, compared to the series LC

Analyzing the Gain- Phase Relationships of the LLC Tank Circuit

Looking at Figure 12, we see that in a series resonant LC, for any load, we always had to be to the right of the resonant peak (~159 Hz), for the phase to be greater than 0° ("inductive", i.e. for current to lag the voltage as desired). That would mean ZVS was possible only for >159Hz (potentially up to very high frequencies). In the LLC tank, we get a "low-frequency phase boost" arising from the lower resonant peak, which in turn is related to L_2, and this makes a wide range of load resistances with ZVS possible *even at lower frequencies*. In fact when we do the full math, we see that we get two resonant peaks, one formed by C and L_1 (as for series resonant tank), and another one formed by C and (L_1+L_2). Hence we have the resonant frequencies designated "f_{RES_1}" and "f_{RES_2}" shown at the bottom of Figure 12. Numerically, for the values chosen, the second peak appears very close to 50Hz. So, we can typically operate to the *right of 50 Hz and up to 159 Hz* for all loads ranging from open-circuit to R = 200 Ω (it seems so far), and achieve ZVS. For loads greater than that (say R < 200 Ω), we have to operate above 159 Hz, otherwise we cannot get ZVS because the phase angle is negative ("capacitive").

You might say: "So what" (if we can operate down to lower frequencies for a range of loads)? That by itself doesn't seem to matter! In fact we usually want to operate at *high* frequencies anyway, and want to try and ensure ZVS at *high* loads primarily. So what is the advantage of the LLC? The real advantage of the second peak in the LLC tank shows up only in Figure 11 (in the *gain* plots). In this figure, we have plotted out the gain, which we often call "conversion ratio" (i.e., output voltage of the LLC stage, divided by input voltage, where output is just the voltage across the inductor in our simple AC case so far), versus frequency (for different loads).

For very high loads (e.g. R = 1), we get a conversion ratio *always less than 1*. That is just like a series resonant LC in which we only get step-down ratios. If we decide to operate in that region with our LLC tank, we can do that. But there is a major stumbling block: For very light loads (such as 1k or 5k, as shown in Figure 11), we cannot even get much less than gain = 1, unless we move to infinite frequencies. We could have a case where we can't even regulate anymore over the full load range. Or worse, maybe we deigned our entire converter to step down, but at light loads the control loop moves in the "wrong direction" --- and ends up on stepping up (perhaps because the shape of the curve, in particular its slope, flipped signs as we moved from one side of the resonant peak, to the other side side. Yes, maybe we can "preload" the converter (at a huge expense to light-load efficiency), so our tank circuit never sees very light loads. In fact, we always need to do that in the SRC just to avoid the huge frequency variation we mentioned previously, but we could end up in the same situation with LLC too, if we are not careful in designing our region of operation, and our control strategy.

Therefore looking at things very practically, we decide we should ***not try and operate an LLC based converter to the right of the higher-frequency resonant peak*** (159 Hz in our case).

Let us keep in mind that designing our system to be "step-up" or "step-down" is actually *our* initial design decision, whatever the relationship of the input and output voltage levels. Because eventually, *we will use a transformer* with an appropriate turns ratio to correct for that. For example, we could design our LLC for a 1:1.1

ratio (an output 10% higher than the input), and then use a transformer with a Primary to Secondary turns ratio, of say 2:1, to bring the output down to about half the input voltage. That is very similar to what we do, or can do, with the conventional Flyback topology too. For example, in a typical universal-input off-line (AC to DC) Flyback power supply, we create an "invisible" *intermediate Primary-side regulated rail* of about 100V. This is for all practical purposes, the output rail as "seen" by the Primary side, and it regulates (and sets duty cycle) to keep this level fixed at 100V. But after that, a 20:1 transformer ratio is inserted via the transformer, to produce 5V from the 100V intermediate rail. Likewise, in *transformer-based resonant topologies* too, we have an additional degree of design freedom coming to us from the turns ratio, and we can voluntarily pick whether we want to operate the Primary side to act as a step-up, or step-down stage. As a corollary, since we will use a transformer anyway (also for creating the magnetic elements of the LLC), isolation is a natural advantage that accrues from an LLC converter (whether we like it/need it, or not).

We have decided not to operate to the right of 159 Hz in Figure 12. What about to the left of the second (lower-frequency) peak: i.e. below 50Hz? Unfortunately, looking at the phase diagrams in Figure 12, we realize that the phase is below 0° in that region, and therefore the network will capacitive to the input source. We know that is not going to lead us to ZVS operation. In other words, finally ***we are left only with the region between 50Hz and 159Hz to operate over the entire line and load variation***. We have no other practical and unambiguous choice really, given the guiding criteria discussed previously and the shape of the gain curves. But one question remains: does this (available) region (fully) meet *all* our needs? Do we need to look more closely?

Let us look again at Figure 11 and see how we intend to implement *line and load regulation*. We can see that, depending on the load, we can get either step-up or step-down (i.e., conversion ratio greater or less than unity). Most of the curves are in fact with a step-up ratio. We can actually close-up in this vital region, as we do in Figure 13, and we realize that the curve of R = 200 Ω, which on the basis of its positive (inductive) phase had been previously deemed "ZVS-capable" and therefore acceptable, won't *work for an entirely different reason*. That is related to one of the two guiding criteria we talked about earlier: do we have a "simple way to implement an unambiguous, almost "knee-jerk" (rapid) response to disturbances, *analogous to what we do in conventional power conversion?"*. For the 200 Ω curve we are failing this requirement for line regulation. We can see from Figure 13 that the 200 Ω curve slopes *downwards* as frequency decreases from f_{RES_1} to f_{RES_2} (159 Hz to 50 Hz). This is *opposite* to all other neighboring curves here. In all the other cases, based on Figure 11 we have obviously designed the converter in such a way that if input falls, we lower the frequency (in knee-jerk fashion), to correct for that. But for the 200 Ω case (as per the closeup in Figure 13), decreasing the frequency will cause the conversion ratio to fall even more, so the output *will collapse*. This is not in the "right direction" for correction, commensurate with neighboring trends and our control strategy.

We therefore conclude that:

Resistive loads less than 200 Ω are definitely considered "overloads" (based on the LLC tank circuit values chosen so far). Output foldback will occur in this case. These values cannot be supported, even though the phase seemed right.

This "dropping off" of gain is coincidentally almost commensurate with the phase just going negative (below zero degrees), as indicated in Figure 13. While plotting these curves, the spreadsheet was set to "test the gain versus the phase condition", and "appear dotted" if the phase was reported to be negative. In general we can safely conclude that whenever the gain starts "heading down", the phase simultaneously is close to dropping below the zero degree line (it becomes negative, and thus capacitive). Therefore, for *two* reasons, not one, the dotted part of the curves in Figure 13 is considered "no man's land" from now on. If the system enters that region, it won't be destroyed, but its efficiency will suddenly worsen (no ZVS), and also, any attempt at output correction, will most likely cause output foldback. On this basis, we rule out 200 Ω --- we can see from Figure 13 that it just doesn't go far enough into the lower frequency region (starting from f_{RES_1}), before it turns "dotted". We also conclude that the max usable load for any LLC converter, based on the LLC tank circuit values we picked in this example, is ***250 Ω*, but also, adding some design margin, *we fix preferably 300 Ω***. Higher loads than that (i.e. smaller load resistance values), are "overloads" by definition.

At very light loads, the frequency will shift close to 159Hz, not more (certainly unlike an SRC which tries to go to infinite frequencies to regulate at light loads).

At very heavy loads, the frequency will shift closer to 50Hz.

We can easily sense, based on the equations for f_{RES_1} and f_{RES_2}, that increasing L_2 significantly will move f_{RES_2} towards much lower frequencies. Though that may help in some way, it is clear that there is a penalty for making L_2 much larger than L_1. For one, the frequency variation spread, going from min load to max load, and from low line to high line, will become much larger.

We have decided we need to ***design for conversion ratios higher than unity*** so we can cover the full load range from 250 Ω (peak load) to higher values (unloaded). In particular, to ensure ZVS, we will set the LLC converter at peak load, at high line, to be at the encircled point at the lower right tip of the shaded area in the LLC phase diagram of Figure 12. At that particular design entry point (shown more clearly and precisely with an ellipse/circle in Figure 13), we will have a frequency closer to the higher resonant peak (~ 125 Hz in our case here), and a corresponding gain (conversion ratio) of about **1.05**. This ***gain target of 1.05 at max load and high line is actually our universal recommended entry point for all LLC converters designs***, because otherwise the gain curve "droops", and the correction loop will take it in the "wrong direction" (foldback) as discussed previously. If we set it closer to 1, say at 1.02, our design entry point will shift closer to the higher-frequency resonant peak. But that is too marginal.

The full optimum region of operation is also shown (shaded) in the gain curves of Figure 13. This area is the part of the gain curve *for which phase angle is positive (inductive)*. As indicated, we have "AND-ed" the gain and phase curves using Mathcad, so the curves become dotted if unacceptable. As also indicated, this does not mean the system will necessarily stay in this optimum region, or we are somehow constraining it too. That locus of movement is determined solely by the regulation loop, as it relates to the specific line or load variations. So certainly, our selected *design entry point must lie within this shaded region (at the circle/ellipse)*, so that we get ZVS at max load and at high line, but we *may just lose the advantage of ZVS if we leave this shaded area.* For example, if we set the max load as 300 Ω, we will be able to reach a conversion ratio of about 1.18, as per Figure 13, before we leave the ZVS region. That seems to equate to an allowed 18% reduction in input (but we will soon show it is only 12% actually). Either way, we can produce ***derating curves*** to show how we can ***tradeoff max load versus input operating range,*** *weighing it against the price of overall expected efficiency, provided of course the range can be achieved or is acceptable in terms of our guiding criteria*.

The fear of predicting and/or really losing efficiency at low line is perhaps one reason LLC converters are still largely being used only where the input is relatively stable (and comfortably so), such as where the LLC converter is driven off the output HVDC rail of a front-end PFC stage. We can see that if we really want wide-input variation, we have to *overdesign* the converter (300 Ω in our example, will only get us 12% input voltage variation factor as explained later). Holdup time can also be a major issue in AC-DC designs in which the LLC converter stage typically follows a conventional PFC stage.

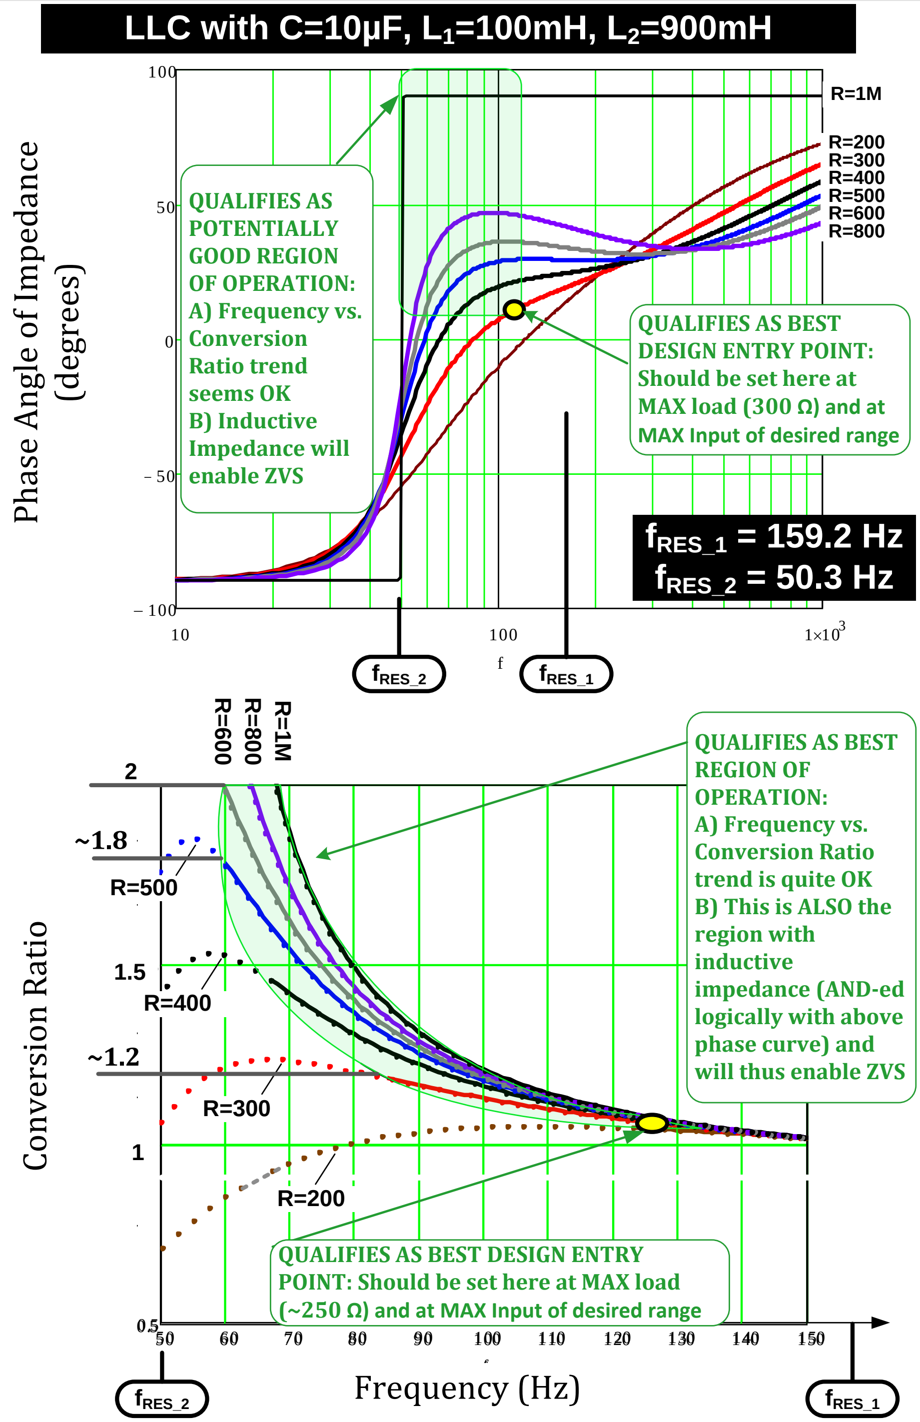

Figure 13: Closeup of phase for different loads, and corresponding conversion ratio (gain) in the usable area of the LLC tank circuit between the two resonant frequencies

Two Resonances in LLC Tank

In related literature, it is often said that the LLC converter has two resonant peaks: one a series resonance between C and L_1, and another "parallel" resonance between C and (L_1 + L_2). So, people call it by many names: f_r, f_s, f_p, f_m, f_o, f_∞ and so on. It can become confusing, which is why we have just preferred to call them f_{RES_1} and f_{RES_2}, based on the simple order they actually appeared in our discussions. Yes, f_{RES_1} does come about from the resonance between C and L_1, whereas the lower frequency peak, f_{RES_2} comes from the resonance between C and (L_1+L_2). But are these series or parallel LC resonances?

We should not get confused by the fact that the conversion ratio (gain) at the first (higher) frequency is less than 1, as is typical of a series LC circuit, whereas the gain at the second (lower) frequency is greater than 1 (which is perceived as true for parallel resonant circuits, in tuned circuits for radio applications etc.). To really find out whether the peaks are series or parallel resonances, we need to plot out the input impedance (as seen by the source). This is presented in Figure 14.

It is interesting that going from a shorted load to no load, the impedance presented to the AC input shifts between what are clearly *two resonant peaks* of series resonant characteristics, because in parallel LC resonance, the impedance becomes very high at resonance, not low as indicated by Figure 14.

One of the contributors to efficiency is switching losses, which we have tried to minimize by invoking ZVS. But that advantage could be easily lost by excessively high conduction losses, as caused by *high circulating currents*. One of the contributors to that circulating current is indeed L_2, which is why typically, L_2 is kept at least 5× larger than L_1 (usually less than 10× though, as explained further below). Especially when we insert a transformer with output diodes, we will see L_2 significantly affects the current distributions in the LLC tank. Therefore, it is recommended to ***keep L_2 between 7 to 11 times L_1 in any practical LLC converter. In our case we have fixed on a factor of 9.***

Note that in a parallel LC, the circulating current sloshes back and forth between the L and C, so the input source may never "see" the high current. But in a series resonant case, the current does pass through the input source and is inversely proportional to the *impedance* the network presents at its input terminals (to the source). If the impedance is high, we will get small circulating currents (and higher efficiency). In Figure 14 we see that there is a *shaded encircled area* of high impedance for almost any load. It is recommended in related literature that we should try to remain here as much as possible. But note that our previously honed choices, R = 250 Ω or R = 300 Ω, fall in a very flat part of the encircled region. In other words, our previous design choice is actually good even from the viewpoint of low conduction losses. We have to do nothing more to optimize.

Just for information, we mention that in related literature, the cusp labeled "f_X" in Figure 14, is mentioned as some sort of LLC converter "design target" because it offers high impedance. Its equation is $f_X = 1/ 2\pi\sqrt{[C\times\{L1+(L2/2)\}]}$. We have however preferred to generalize our approach by describing a certain load resistor corresponding to max load at high line, one which a) gives an unambiguous direction of correction, and b), makes the network appear just slightly inductive to the source (for ZVS). *The desire for high input impedance is automatically taken care of* as we see from Figure 14.

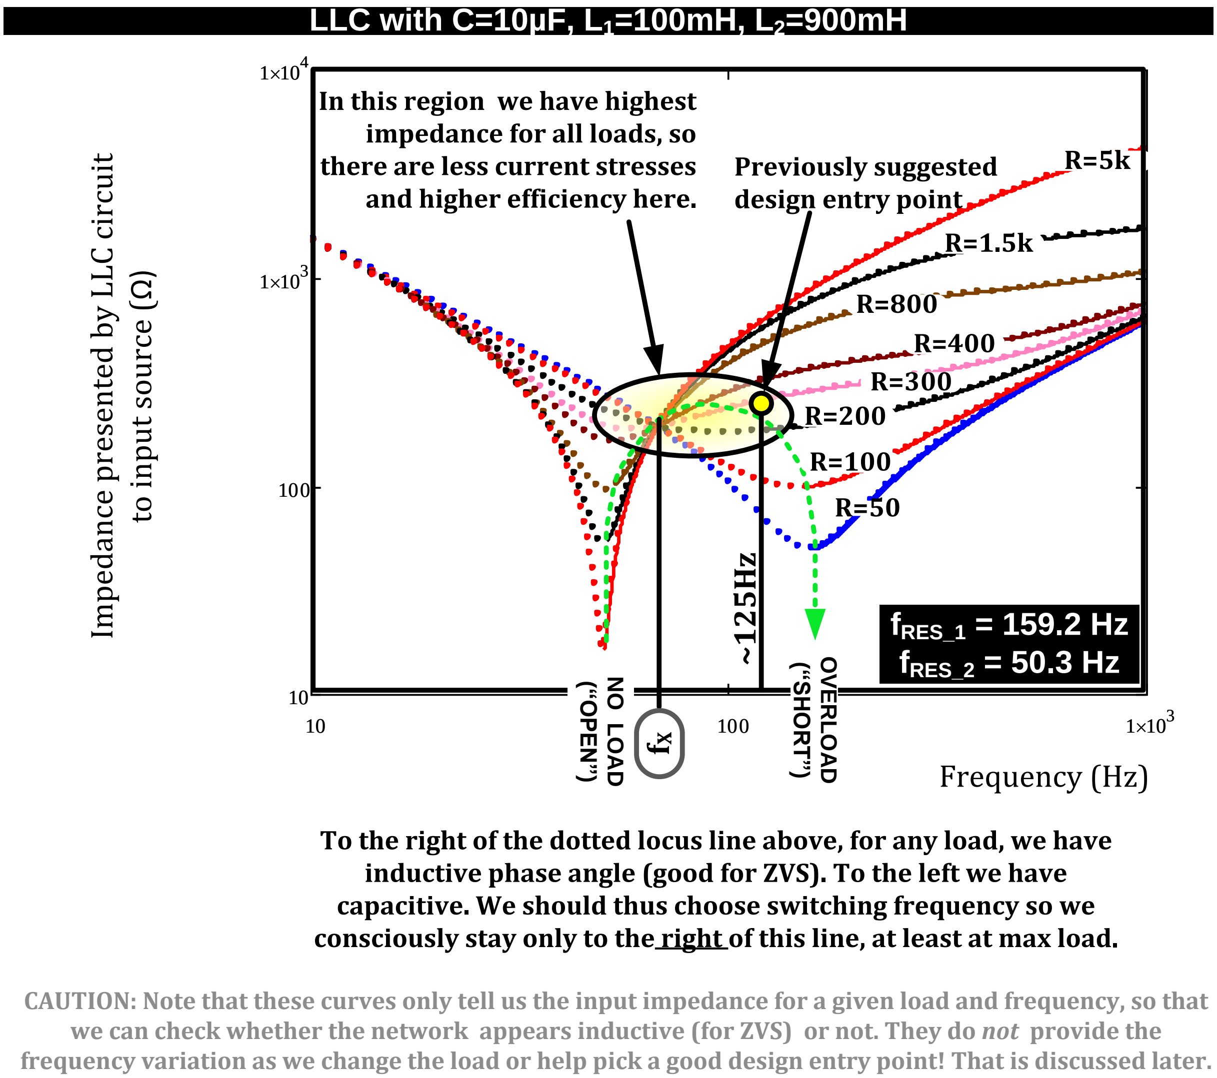

Figure 14: Plotting input impedance of the LLC tank circuit

PART 3: Microsemi Innovation

One of the perceived disadvantages of the LLC topology is that it handles input line variations poorly. This has restricted it into a niche market, mainly AC-DC, and placed after a Power Factor Correction stage supplying a steady 400V DC to the LLC Half-Bridge. It certainly could be of great advantage to achieve a wide input variation with LLC, say 2:1 or better.

Another perceived disadvantage is that at light loads, LLC engineers are still not able to guarantee operation is restricted at the upper end, especially at light loads. So most LLC controllers only place bounds on the lowermost frequency. However it is advantageous to specify the upper limit too, because for example, the designer may wish to stay decisively below 150kHz, which is the start of the CISPR 22 (EN55022) conducted emissions band. If operation can be guaranteed below 150 kHz for a wide line variation, say 4:1 (e.g. 100VDC to 400VDC, or 110VAC to 240VAC), along with a full load variation (zero to max), it can be of great help in almost eliminating input filtering, especially in AC-DC applications. It could also be of great help in minimizing effects of switching converters on signal integrity where power and data coexist, such as in networking applications.

However the biggest stumbling block to widespread adoption of LLC in industry is simply this: it does not yet fill into a predictable design schedule, so essential in modern development…and that in turn is because engineers still adopt a trial and error approach to fixing the L, L and C values, along with the turns ratio.

In December 2012, Microsemi achieved a breakthrough, and as a result, a predictable design methodology is emerging to account for all the weaknesses above --- wide input variation (4:1 estimated), set upper and lower frequency limits and a commercially acceptable design procedure with minimum bench tweaking.