An Ultracompact (9.2 W/cm^3), Electrolytic Capacitor-free, Single-Phase, Non-isolated 2 kW Photovoltaic Inverter

Alexander Huenten, Team: AHED, Registration Code: 545-UMNYeK-59906,
Duesseldorf, Germany
alexander.huenten@ah-ed.de

***Abstract*—This technical approach document describes at a high level the strategy and key innovations used to address Google's Little Box Challenge**

***Keywords*—Single-phase inverter, active ripple filter, GAN FETs, ZVS, interleaving, phase shedding, MLCC, FPGA, soft-core processor**

I. Introduction

A high level block diagram of the inverter achieving a power density of 9.2 W/cm^3 (150.78 W/in^3) in a volume of 217.4 cm^3 and weighing just 300 g is shown in **Fig.1**. Several uF of input capacitance (C_DC) provide HF-decoupling for the switching cells. The HF-Inverter does create a rectified sinusoidal with 340 Vpk from the DC input voltage of 450 V by 4 interleaved phases of the basic switching cell **(Fig.2)**, which is later elaborated on. The active Ripple Filter (RF) does absorb or inject energy on the DC input in order to suppress the 120 Hz power ripple created by the inverter. It consists of 2 interleaved phases of the same basic cell. The LF-Inverter does transform the rectified sinusoidal into a true sinusoidal and is implemented as a simple full bridge since this allows for a straight forward implementation and no HF common mode noise is injected. There are no strict requirements on the switching behavior of the semiconductor devices since they are only switched at 60 Hz. Alternatively, a single 60 Hz-switched half-bridge ("Slow Leg") in addition to the 4 interleaved "Fast Legs" of the HF-inverter similar to the setup of a bridgeless totem-pole PFC could be used. This would reduce the semiconductors in series in the power path from 3 to 2 at the cost of a more complex control scheme and a higher common mode noise at the zero crossings, which makes this solution less attractive for photovoltaics with relatively high to-ground capacitance of the solar panel as compared to e.g. fuel cell applications. The EMI filter is supposed to attenuate conducted switching noise to a level compatible with FCC Part 15 B. The inductors used for the basic cells are on the order of µH, the input/output capacitors on the order of uF, the switching does occur into the MHz range.

II. Basic Switching Cell and ZVS

The basic switching cell used in this design is shown in **Fig.2**. At first sight, it's simply a synchronous buck or boost cell (depending on looking from which side). Usually, such a

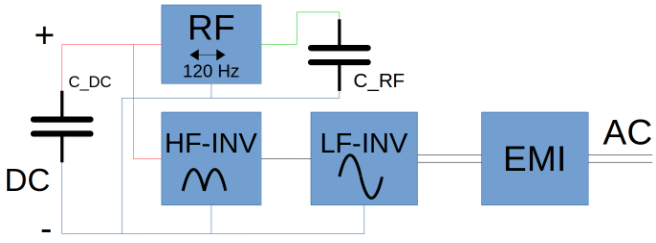

Fig. 1. High-Level Block Diagram

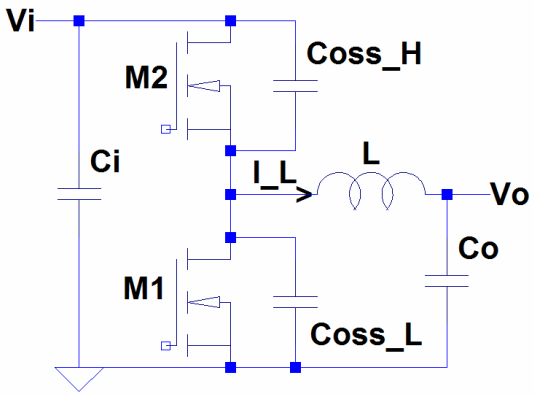

Fig. 2. Basic Switching Cell

cell is operated with a plain PWM modulation that entails a hard switched transition for at least either high-to-low or low-to-high transition at a given operating point. At elevated switching frequencies and especially for high voltage Super Junction FETs this does cause significant losses due to cross conduction, Eoss/Qoss and for SI-FETs reverse recovery of the body diodes. In this design, however, the cell is operated with resonant transitions only. This is achieved by making sure that before a transition, which is initiated by turning off the previously on FET while the partner FET stays also off, an appropriate amount of energy is stored in L to create the wanted resonance with Coss_L/Coss_H (for positive/negative I_L the inductor will lower/raise the switching node). The appropriate amount of energy is determined on-the-fly based on Vi, Vo and the Coss(V_DS) and Qoss(V_DS) profiles of M1 and M2.

If the turn-off occurs sufficiently fast, then there is almost no turn-off loss thanks to the highly non-linear Coss of typical high voltage FETs which acts as a turn-off snubber and does effectively allow for Zero Voltage Switching (ZVS) turn-off.

Eoss is ideally entirely recovered, however practically there are losses associated with Eoss due to ohmic losses and core losses in the partner inductor L.

The time it takes for the resonant transition to complete (both FETs are off) is precisely calculated so that the FET to which the current commutes is turned on as early as possible but as late as needed for ZVS turn-on. This does effectively eliminate any significant body diode conduction with the associated conduction losses and body diode recovery losses (especially for SI FETs).

The results of this are extremely low loss transitions. For SI and GAN FETs selected for this project as low as 1.25uJ are consumed per transition and FET (at MHz scale frequencies) for a ~500W cell. This does include all losses of the power stage related to the resonant transitions. Comparing SI to GAN FETs, GAN FETs have an advantage due to their smaller Coss, Qoss and Q_g to R_DSon ratio. Most relevant for this design is that SI FETs lead to physically larger inductors due to the larger energy requirement for the resonant transitions and consequently higher stress on them which would increase the total volume by ~ 30%. Therefore, GAN FETs had to be used for the final design.

Another important advantage of these resonant/soft transitions is that they entail less HF switching noise as compared to hard transitions.

Not surprisingly, these advantages come at a cost: The maximum instantaneous frequency is implicitly set by the operating point (Vin, Vout, Iout) and the constraint for resonant transitions. Under the further constraint of minimum solution size there is a narrow band of feasible operating frequencies for a given operating point. For a circuit with widely varying Vin/Vout ratio like a sinusoidal inverter the operating frequency does vary likewise and is constantly modulated. This does entail a significantly more complex control than a simple PWM modulation; especially if several phases have to be synchronized and should be switched on/off dynamically.

The high frequency ripple current in the inductor required for the resonant transitions does put high stress on the inductor (core loss, HF winding loss) so that it has to be carefully designed. This ripple is seen at the output/input capacitor and consequently with a certain attenuation at the output/input itself, as well, so that usually several interleaved phases are required to mitigate this. When doing so the effective ripple frequency is furthermore increased which eases post filtering.

For given FET and inductor technologies the maximum feasible frequency/minimum solution size is limited by the fact that if the losses/volume associated to the HF ripple current approaches the losses/volume associated to the load current a further increase in frequency won't pay.

In fact, the general concept of resonant transitions, already 4x interleaved, is by far no new concept and was practically used at least as early as 1992 [1]. However, it becomes more and more attractive with higher focus on energy efficiency, solution size and the availability of high computational power, but low electrical power and low cost digital signal processing.

III. DIGITAL CONTROL

The control of this device is implemented entirely in the digital domain. The input signals for this control are only input/output voltages and synchronization signals from each switching cell. It's implemented solely in an FPGA device; no dedicated processors are used. Within the FPGA the control occurs at 3 levels of hierarchy:

At the highest level is a proprietary floating point soft-core processor "CC_MAIN". It is responsible for on-the-fly calculation of all aspects of the resonant switching and all lower speed (~50 kHz) PID regulators. It has HW accelerators for more complex mathematical operations like DIV, SQRT, SIN, COS.

At the next level is a proprietary fixed point soft-core processor "CC_ILC". It is responsible for feeding the modulator engines with data, based on what is calculated by "CC_MAIN". Moreover, it manages a modified modulation scheme around the zero crossings for the HF-Inverter since the feed-forward assumptions (e.g. constant Vout during the switching cycle) made by "CC_MAIN" are not adequate here anymore.

At the lowest level are modulator engines for each switching cell. They create gate control signals on an ns scale and are responsible for phase synchronization. Due to the constantly varying switching frequency a high speed phase correction loop is required to keep the individual phases at their scheduled phase offsets. This phase correction is performed for each slave phase up to twice per switching period so effectively at MHz scale. The modulators due to their supervision of the phase synchronization signals do effectively represent cycle-by-cycle current limiters that can detect severe errors on a µs time scale.

A further noteworthy aspect is that dynamic phase shedding does occur. Within a line half period the number of active phases can be varied from 1 to 4 for the inverter and 1 to 2 for the ripple filter under the constraints of highest efficiency and lowest ripple.

The overall FPGA design does only consume 5000 LUTs, 2500 FFs, 50 kByte of BRAM and 6 MULs in a Spartan 6 device.

IV. INPUT RIPPLE

Using the traditional large DC input capacitor would require ~1550 uF of ideal capacitance to fulfill the competition requirement[1]. Using commercially available 450 V aluminum electrolytic capacitors they alone would consume roughly the volume of the entire device presented here. The root cause is that only a small fraction (+-2.46 J) of the energy storage

[1] See Appendix A for a derivation

capability of the capacitor (157 J) is actually used due to the small allowed voltage excursion.

Theoretical solutions to this issue are obvious: Either a battery-like storage device has to be used (a) where most of the energy is stored within a small voltage band or b) the voltage excursion seen by the capacitor has to be increased.

a) Commercially available batteries are not capable of handling charging/discharging cycles at 120 Hz at an attractive power/energy density.

b) The power density of commercially available electrolytic capacitors is limited by their high ESR so that they are not attractive in this case. Commercially available supercapacitors suffer from their relatively large ESR at 120 Hz, as well, although graphene double-layer capacitors are a promising upcoming technology [2]. Film capacitors have a very high power density at 120 Hz but their energy density is not attractive. The viable alternative left are MLCCs that provide high power and acceptable energy density at 120 Hz. MLCCs which tend to be physically small as individual devices and need to be connected in parallel have the added advantage that they can be arranged in such a way as to make optimum use of (otherwise possibly unused) volume.

As indicated in **Fig.1** an array of high voltage MLCC capacitors (C_RF) forms the energy storage of the Ripple Filter and can be operated at an energy usage ratio

$$\Delta E / E_{\max} = (E_{\max} - E_{\min}) / E_{\max}. \quad (1)$$

of ~ 50 % at full load.

It should be further noted that this Ripple Filter, within its energy storage and power limits, can principally be controlled to completely eliminate the input ripple for an arbitrary output current shape; e.g. higher order harmonics could be cancelled, as well.

V. MINITURIZATION

Under the constraint of soft-switching, which is considered essential to achieve highest power density as elaborated above, the switching frequency is increased as much as reasonable for the available device technologies (FETs, inductors). For ultimate power density GAN FETs have to be used due to their superior FOMs for Coss, Qoss and Q_g to R_DSon.

The soft switching method used here entails high inductor ripple current at frequencies into the MHz range. Only core materials (ferrite) with lowest loss figures can be used. The winding geometry is optimized to keep the AC winding loss in a reasonable range. The core shape ("planar magnetics") is selected as to provide a large surface area to volume ratio in order to allow for operating the material at high core loss density and thereby not unnecessarily far from its saturation limit.

The all digital control scheme, centralized in a single FPGA device, allows for operating the subcomponents at optimum operating points for varying load conditions, e.g. through techniques like dynamic phase shedding and precise interleaving. Interleaving does significantly mitigate the high input and output ripple currents associated with the soft-switching scheme, increases the effective ripple frequency seen and entails smaller basic cells, which lend themselves to being combined to a compact device.

Components are selected or designed under the constraint of allowing compact functional groups at regular geometric size to not keep volume unused when combining them to the complete device.

When striving for very high power density the effect of auxiliary components must not be underestimated. A commercial auxiliary supply 450 V to 12 V, 5W (for gate drivers …) could consume a significant portion of the overall volume and might not integrate well with other functional groups due to its geometry. Therefore, another downscaled instance of the basic cell in **Fig. 2** is used to generate the main auxiliary rail. At initial power up a 2nd mini FPGA will control this stage with a simple hysteresis scheme. Once the main FPGA has been powered up it takes over control, with a ZVS scheme. Without excessive optimization effort the power stage (ignoring the FPGAs) does consume only ~ 3.2 cm^3. Thanks to ZVS it generates low noise as e.g. compared to typically hard switched fly-back topologies.

The synchronization circuitries for the ZVS as used here require several isolated auxiliary supplies (3.3V, < 100 mW). In order to spend as little volume as possible on these, a multi-MHz isolated push-pull auxiliary supply was designed. It's again driven by the main FPGA (fully feed-forward) and does consume only ~ 0.12 cm^3 (8 x 6 x 2.5 mm) per instance (ignoring the FPGA).

VI. THERMAL MANAGEMENT

Designing for ultimate power density does inevitable lead to a thermally limited design since smaller does entail higher loss density and less surface area to dissipate it, which entails more afford (power, volume) to keep temperatures within limits.

One strategy for systems with a single/a few hot-spot(s) can be observed when looking at a modern powerful laptop. The cooling system consisting of copper heat sink and fan having a size of ~ 8 x 8 x 1.5 cm can easily dissipate 50 W generated by CPU and/or GPU.

This technique was not considered appropriate for this design since the initial power loss model showed that power loss would be distributed across many components, including several interleaved half bridges, the inductors, the LF inverter FETs and the EMI filter. Attaching all these components to volume efficient heat sink/fan structures, especially when also considering electrical constraints was not considered realistic. Another issue with heat sinks is the danger of increasing parasitic capacitance of switching nodes, which would entail increased EMI and loss. Instead of designing the system around a heat sink the system was designed to have an inner structure similar to a heat sink. The individual functional groups are designed to form slices with a thickness of only a few mm to ~ 10 mm. These slices are stacked at distances of 1 to 2 mm and thereby resemble fins of a heat sink. The PCBs in those slices conduct heat from local heat sources like the FETs

and dissipate it via their surface or even other components on the same PCB. Planar magnetic components which are a major source of heat, as well, expose their large surface area to the air stream in between the slices, which allows them to be operated significantly above usual core loss density recommendations without excessive temperature rise. In **Fig.3** this concept is shown with a subset of the components of the overall device.

A standard 50 mm fan, which does effectively cover the entire front of the device, pushes air through the device. Thus, about 75 W of thermal power at full load can be dissipated without any dedicated heat sink structure. The actual outer surface of the device has only a minor role in dissipating heat; the much larger inner surface with the help of forced convection by the fan does dissipate the generated heat.

VII. EMC

The ZVS scheme was among other aspects selected for its favorable HF switching noise[2] to switching speed/low loss behavior. The principal drawback of higher ripple current of a single phase for the used ZVS scheme is mitigated by interleaving phases, which does relax the requirements on DM filtering. The lack of any heat sink, which would tend to increase CM noise due to the additional switching node to power earth capacitance, relaxes the CM filter requirements. The fact that most of the time switching does occur in the MHz range does obviously reduce the filter size. The varying switching frequency does behave similar to a strong spread spectrum technique and does distribute the switching noise across a wide frequency band. Nevertheless it leads to pretty low frequencies around the zero crossings which does require a noteworthy DM rejection here, as well.

The conducted EMI filter has to fit into the slice topology with minimum additional volume. Under this constraint especially the magnetic components for DM and CM rejection are optimized considering material, shape, topology (single vs. multiple stage) for the noise spectrum generated by the inverter. A special challenge is to limit crosstalk from other physically close functional groups into the final stage(s) of the EMI filter. Promising concepts for integrated and/or active EMI filters could not be further investigated due to a lack of time.

The peak slew rates achievable with GAN FETs and similarly with certain SI FETs can result in ringing in the 100 MHz range and pose a noteworthy risk of radiation. The most obvious remedy to reduce this is to minimize the loop area of such currents and select components for smallest parasitics (especially lead inductance). The most critical case is the loop formed by Ci, M1, and M2 in **Fig.2**. The next critical loop is the one formed by the gate driver through the FET's gate and source, which is not explicitly shown. Under the constraint of discrete devices these loops were reduced to an absolute minimum. Several product announcements like the

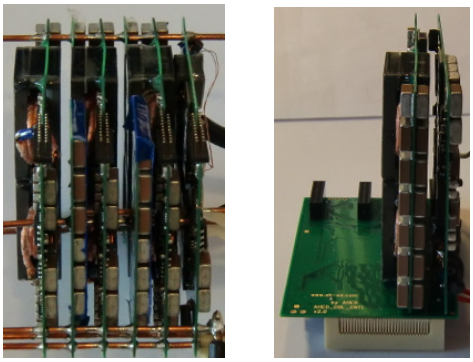

Fig. 3. Left: Several connected power stage slices, Right: Two power stage slices plugged into the control board

"LMG5200" by TI[3] indicate that within the next years further improvement will be possible by integration of the two half bridge FETs with isolated gate drivers and preferably Ci in a single package. Unfortunately, such developments are currently beyond the resources of AHED.

VIII. COST EFFICIENCY

This device has no strict requirement for GAN FETs. With a moderate decrease in power density SI FETs can be used. No expensive magnetic materials are used, just ferrite suitable for high frequencies and iron powder. The windings for the planar cores could be realized as PCBs as commonly done today. No high speed analog circuitry is needed, just comparators and relatively low speed ADCs for voltage sensing. The design is scalable; phases can be easily added or removed for different power levels. Due to the maximum thickness of ~10 mm of the slices the device can be reorganized from a cube to a flat plate that might even be passively cooled if attached to a structure with appropriate thermal resistance. Even though FPGA technology is used the highly resource optimized design does allow for using sub 10$ parts. Almost only SMD technology is used which results in low assembly cost. Only FR4 PCB material and standard reflow soldering is required. The distributed power loss without extreme hot spots does not necessitate dedicated heat sinks. Today just the MLCCs required for the Ripple Filter pose a note-worthy cost burden of ~60$/kW. However, the prices of MLCCs are constantly falling and there are promising upcoming alternative technologies like graphene ultra-capacitors.

[1] D. M. Sable, F. C. Lee and B. H. Cho, "A zero-voltage-switching bidirectional battery charger/discharger for the NASA EOS Satellite", Applied Power Electronics Conference and Exposition, 1992. APEC '92. Conference Proceedings 1992., Seventh Annual

[2] J. R. Miller et al., "Graphene Double-Layer Capacitor with ac Line-Filtering Performance", Science 329, 1637 (2010)

[2] Noise related to the transitions/ transition speed (e.g. ringing) and not the switching cycle frequency

[3] 80V GAN Half Bridge Power Stage with integrated gate driver

Alexander Huenten

Academic background:
Mr. Huenten did study Electrical Engineering and Information Technology at RWTH Aachen from 2005 to 2009. He graduated with honor as Dipl.-Ing. in 2009 with the final thesis “Development of a SDR Receiver for Radar based Detection and Tracking of Objects in Space”. During his studies he received a scholarship from the “Studienstiftung des Deutschen Volkes”. Contacts made at a meeting of the “Studienstiftung des Deutschen Volkes” in Aachen allowed him to develop FPGA based control systems for large scale LED video walls for a company in Duesseldorf while still studying.

Professional background:
Since graduation Mr. Huenten is working as a freelance developer under the label AHED - www.ah-ed.com - with focus on systems around FPGAs. Areas of expertise are LED lighting, industrial control (from PLC over motor control to high frequency, high power inverters for ultrasonic applications, including the digital control section as well as power electronics), image/video processing.

ABBREVIATIONS

BRAM: Block RAM, FPGA resource
Bridgeless totem-pole PFC: Totem-pole PFC without an input diode bridge
CM: Common Mode
Coss: Output capacitance of a FET
DM: Differential Mode
EMI: ElectroMagnetic Interference
Eoss: Energy stored in output capacitance Coss of a FET
Fast Leg: Half bridge switched at tens/hundreds of kHz
FET: Field Effect Transistor
FF: Flip Flop, FPGA resource
FOM: Figure Of Merit
FPGA: Field Programmable Gate Array
GAN FET: Gallium Nitride FET
HF-INV: High Frequency INVerter
LF-INV: Low Frequency INVerter
LUT: Look Up Table, FPGA resource
MLCC: Multi-Layer Ceramic Capacitor
MUL: Multiplier, FPGA resource
PFC: Power Factor Correction
PWM: Pulse Width Modulation
Qoss: Charge stored in the output capacitance of a FET
Q_g: Gate charge of a FET
RF: Ripple Filter
R_DSon: Resistance between Drain and Source terminal of a FET when turned on
SI FET: Silicon FET
Slow Leg: Half bridge switched at grid frequency
Totem-pole PFC: PFC using a half bridge as the main switching cell as opposed to the usual low side MOSFET, high side diode pair
V_DS: Voltage from drain to source of a FET
ZVS: Zero Voltage Switching

APPENDIX

A. Derivation of an estimate of the required DC input bulk capacitance without active Ripple Filter

Operating point with worst case relative input ripple current and voltage is: S_out = 2 kVA @ |PF| = 0.7
-> Average output power: P_out = 1.4 kW
-> 120 Hz ripple power waveform at the output: P_out_ripple(t) = 2 kVA * cos(2*pi*120*t +- 0.7954)
-> Resulting ripple power waveform at the input without any filtering: P_in_ripple(t) ~= P_out_ripple(t)
-> Average input voltage assuming small residual ripple (Rin = 10 Ohm, eta = 96 %):
V_in_DC ~= 450 V – Rin * (P_out / eta) / V_IN_DC ~= 414.85 V
-> Average input current assuming small residual ripple: I_in_DC ~= 3.515 A

Maximum allowed input ripple current (k = 20 %): I_in_AC_max = k * I_in_DC = 0.703 App
-> Resulting maximum input ripple voltage: V_in_AC_max = Rin * I_in_AC_max = 7.03 Vpp
-> The ripple power that needs to be absorbed by the input capacitor is the overall ripple power due to the inverter operation minus the residual allowed ripple power at the input (assuming small residual ripple voltage): P_Cin_ripple (t) ~= P_in_ripple(t) - V_IN_DC * (I_IN_AC_max / 2) * cos(2*pi*120*t +- 0.7954)
= 1854.2 kVA * (2*pi*120*t +- 0.7954)
-> the input capacitor may only show a ripple of 7.03 Vpp for an energy excursion of +-2.459 J around 414.85V DC
-> ~1550 uF
This is based on the original specification of peak-to-peak input ripple < 20% of average input current without restriction to purely resistive loads as added on June 25, 2015. For the modified specification the required capacitance is obviously less.