ON Semiconductor®

Evaluation kit for power line communication — user manual PREVIEW

Pierre Woestyn, Andrey Vangev, Peter Cox et al.
ON Semiconductor Energy Management Products (EM)

2017-06-06

Introduction

This manual describes the evaluation kits (EVK) for the power line communication (PLC) modems from ON Semiconductor.

In addition, we provide some information on how to proceed to an application design (part II).

The information in this manual is focused on the NCN49597 and NCN49599 modems with the ON-PL110 firmware.

To get started immediately with a new evaluation kit, first read [WVC14] and section 1. Then install the required driver and software on a computer (sections 3.1 and 3.2); connect two evaluation kits to the mains and to the computer; and start configuring the modems (section 3.3).

Contents

Part I
Evaluation kit manual

The performance of power line communication (PLC) strongly depends on the environment. Testing in the real world at an early stage is therefore essential. To support customers looking for a PLC solution an evaluation kit (EVK) has been developed.

The evaluation kit allows the user to set up communication between two modems over the power line under control of a PC.

In addition to performance evaluation, the EVK enables an early start of user software development.

The design of the boards is also a good starting point for an application design.

A standard evaluation kit contains two enclosures (each with a motherboard and a daughterboard), two USB cables, a «Getting started» guide, a test report, motherboard and daughterboard schematic, and this manual.

1 Safety

By necessity a large part of a PLC modem board is directly (without isolation) connected to the mains. Safety must therefore be considered carefully.

The main safety risk is electrocution. Alternating current as low as 30 mA can cause heart filibration and death [IEC05]. Under «optimal» conditions these currents can result from voltages as low as 50 V.

An additional risk is posed by the high energy stored in the primary-side power supply capacitor. This large capacitor is charged to the peak voltage of the mains. An uncontrolled discharge will release substantial energy, possibly resulting in injury or damage. A discharge is easily triggered by a moment of inattentiveness with a screwdriver or oscilloscope probe.

When damaged, the capacitor could also explode — a risk inherent in all electrolytical capacitors, but in this case a risk with greater consequences due to the greater stored energy.

Under the recommended operating conditions the enclosure protects sufficiently against these risks.

The evaluation kit is designed to be used in a dry and non-condensing environment.

It should be connected to a normal domestics power socket (measurement category CAT II[1]). If tests must be done linking the evaluation kit to parts of the fixed

[1] Measurement categories were previously called «overvoltage categories» and are still frequently refered to as such. In the latest editions of the IEC/EN 61010-031 standard (since 2002), they denote the energy that is available in case of a short-circuit. CAT I devices are isolated from the mains (cars, battery-powered systems, &c.). CAT II objects are connected to the mains through normal domestic power sockets. CAT III refers to the electrical installation inside buildings; CAT IV to the installation supply sources such as the secondary side of MV-to-LV transformers.

electrical installation such as the fusebox (CAT III), much more energy will be released by incidents such as short-circuits. An additional enclosure is required as a precaution. Contact your sales representative for more information. The evaluation kit should never be used in the supply source part of the mains (CAT IV). Locations such as the secondary side medium voltage (MV) transformers and overhead lines can deliver tremendous amounts of energy and pose high risks.

For 230 VAC systems, only single-phase operation is supported, i.e. with the power cable connected from neutral to phase. Operation across two phases is not possible.

If damage is suspected, stop using the evaluation kit. Contact ON Semiconductor to have it re-tested and repaired.

Your sales representative or field application engineer can also help you with other safety questions — do not hesitate to contact her.

1.1 Primary-side modifications and measurements

From a safety point of view it is of course prefered to use the evaluation kit with the mains-connected part unchanged. However, testing customer-specific parts is often desired. Preferably, the protective cover should only be removed for the modification work (with the power cable unplugged!).

Before starting, ensure the board is safe. During operation some primary-side capacitors are charged up to 350 V_{DC}. Although bleeding resistors bridge all high-voltage capacitors, measuring the residual voltage on critical capacitors remains a good habit.

Following modifications, most measurements can proceed with the enclosure back in place and screwed-down.

If measurements on the primary side are truly needed and the enclosure blocks access, a risk analysis must be performed. Formal safety systems greatly help with this, though this is outside the scope of this document; only a few hints are noted below.

As an obvious safety precaution, shield as many components as possible.

Limiting the energy available during an accident should also be a priority.

Do not rely on the circuit breakers of the electrical installation. Connecting the evaluation kit directly to the mains, while fine for normal use, should *never* be done with modified boards. Instead, use an AC power supply with low trip current setting or an isolation transformer. Isolating the board also protects measurement instruments. It may also be required to avoid tripping residual current breakers.

Consider the required measurement method and instruments carefully to avoid unsafe working and damage. A good guide is available from Tektronix [Tek05].

Never remove the protective earth of a measurement instrument; while this allows floating measurments to some extent, it exposes the operator to lethal voltages on the instrument connectors.

Ensure the voltage rating of oscilloscope probes is appropriate in order to protect the operator and the oscilloscope. High-voltage probes are preferable to general-purpose probes.

Consider slowly ramping up the AC voltage; gross defects such as short-circuits or inverted diodes will be equally obvious at lower voltages but will cause less damage.

Electrical shocks can cause cramps and unconsciousness. Therefore, never work alone on energised low-voltage systems; always ensure a knowledgeable colleague can raise the alarm.

When the fuse fitted on the motherboard has blown, replace it with a suitable type. A fast fuse of 1.6 A or less is recommended.[2]

2 Evaluation boards

The evaluation circuit is split up in a motherboard and a daughterboard (figure 1). Refer to [ON 13, ON 12, ON 14c] for the full schematics.

The motherboard contains the power supply, mains coupling circuit (including the zero-crossing detection) and an USB-to-teletype converter.

The daughterboard contains the modem itself, the power amplifier, the receive and the transmit filter and the protection circuit. Refer to section 19 for a list of available daughterboards.[3]

[2] The high current rating does not stem from the power supply, which consumes 150 mA at most, but from the PLC signal.

[3] If no daughterboard is available for the product or frequency band you are interested in, contact your sales representative to obtain tuned daughterboards.

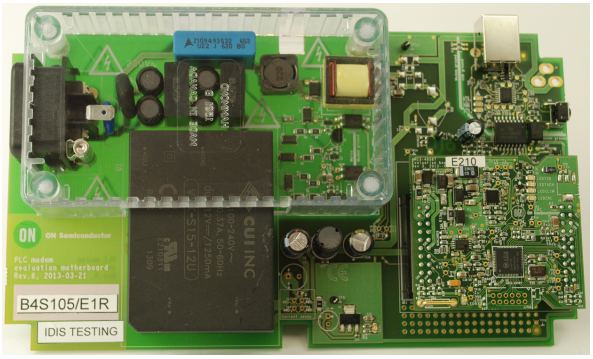

Figure 1: Evaluation boards: motherboard and plugged-on daughterboard.

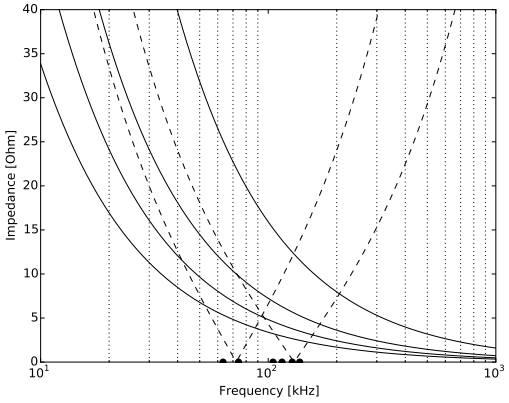

Figure 2: Impedance of typical coupling circuits: 100, 220, 330 and 470 nF capacitors (solid lines, top to bottom); 220 nF+22 µH and 150 nF+10 µH (dashed, left resp. right).

2.1 Power supply

An enclosed switching mode power supply (SMPS) converts the AC voltage to a 12 V, 1 A DC voltage. An input filter is required to comply with EME regulations. In additional, the impedance on the line must be sufficiently high for the PLC frequencies. This precludes the use of a capacitor directly between the phases. Instead, a fourth order LC filter is used. For more information refer to the application note «PLC modem power supplies».

The input range of the power supply is 85–264 V_{AC} and 120–370 V_{DC}.[4] The coupling stage should work with any power line with a voltage of less than 260 V_{AC} (frequency 60 Hz or less) and 370 V_{DC}.

Operation with a power line voltage above the given range is not safe.

Operation with a voltage below the given range is possible, provided the AC-DC power supply is isolated. This may be done by desoldering the first inductor pair of the power supply filter. The board must then be fed from an external 12 V_{DC} power supply through J1.[5]

The 12 V supply is used directly by the line driver during transmission.

A 3.3 V low-drop out voltage (LDO) linear regulator supplies the modem itself, the digital circuits and the additional opamp(s) on the daughterboard.

[4]Consult the manufacturers datasheet for details: for the Mean Well supplies refer to [Mea12]; for the CUI supplies to [CUI13].

The ON-PL110 firmware can operate without zero crossing, if it is configured to do so (page 12).

[5]A footprint is provided but not fitted. Farnell 1737246 or equivalent may be used.

2.2 Zero crossing detection

The low-cost zero-crossing detection circuit proposed in [WVC15] is included on the motherboard. In addition, revision 6 motherboards and later also provide footprints to populate the low-power zerocrossing circuit.

Typically, a delay of 35±5 µs with a jitter of 1 µs is measured between the actual zerocrossing signal and the digital output of the optocoupler.

2.3 Mains coupling

The motherboard contains the coupling circuit with the mains: high-voltage coupling capacitor, inductor (optional) and transformer.

The coupling capacitor is required to block the 50 or 60 Hz mains voltage. Values of 100–470 nF are typical. The impedance of the capacitor near PLC frequencies is shown in figure 2, solid lines.

The IDIS carrier frequencies (63.3 and 74 kHz) and the PL110 dual channel carrier frequencies (105, 115, 127.5 and 137.5 kHz) are also marked at the bottom of the graph.

Often, a lower impedance is desired at the carrier frequencies to allow more power to be coupled onto the mains. To this end a resonant coupling circuit (inductor and capacitor in series) may be used. Figure 2 shows the impedance of two coupling circuits:

220 nF+22 µH and 150 nF+10 µH.[6]

The drawback of a resonant coupling circuit is apparent in figure 2: away from the resonance frequency, the impedance is often higher than with a simple capacitor.

The coupling circuit fitted on the evaluation board depends on the frequency band. For motherboards designed for the IDIS carrier frequencies, a 220 nF capacitor is used with an optional inductor of 10 µH[7]. For ON-PL110 motherboards, 330 nF is fitted.

The capacitor and inductor can be changed easily. Any film capacitor with a suitable voltage rating and a lead spacing of 15[8] or 22.5 mm is suitable. The Epcos B32923[9] series or the Vishay MKT series is suggested. For the inductor, a sufficiently high saturation current must be selected to avoid distortion of the output signal. The Bourns SRR1260 series is suggested.[10]

Four suitable transformers are shown in table 1. Note that the PT10 transformer is not recommended for use with large signals, as the core saturates earlier than the other suggested models.

A footprint is provided for these transformers on the motherboard.[11] By default, only TF1 (ON-PLC-1 transformer) is populated. If you would like to evaluate one of the other transformers, contact your ON Semiconductor sales representative before ordering an evaluation kit.

The daughterboard contains the recommended protection circuit (refer to section 7.4 for details).[12]

As noted in section 2.1, the EVK is designed for a DC or low-frequency AC line with a voltage of 85–264 V_{AC} and 120–370 V_{DC}. For lower voltages, the coupling circuit is still usable, but it is recommended to disconnect the power supply from the mains input. For more optimal designs for DC lines, refer to section 7.3.

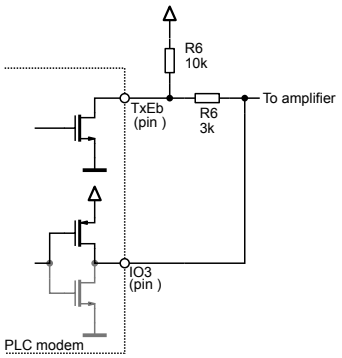

Figure 3: TxEb control circuit. The transistor drawn in grey is not used.

2.4 Line driver enable signal

A standard-compliant KNX PL110 stack needs to be able to transmit half-words. However, the NCN49597/599 can only transmit full bytes. As a workaround, a GPIO can override the TxEb signal. In the evaluation daughterboards this is made possible through the addition of R36 (figure 3). During normal operation and at the beginning of a frame, IO is left in a high-impedance state. To stop the transmission IO3 is set high by the ON Semiconductor ON-PL110 embedded software stack.

2.5 Serial communication

Two UART signals, RxD and TxD, are provided to allow serial communication with the modem. In addition, a *transmission request* (TREQ, on IO2) signal is provided to enable communication with the ROM boot loader of the modem and to enable uploading new embedded software over the UART interface.

Evaluation motherboards also allow triggering a hardware reset over the USB connection. Reading back the error signals of the line driver is also possible.

To avoid parasitic communication paths and protect the computer connected to the evaluation board, the serial communication ground is galvanically isolated from the modem ground.

[6] In both cases, a 75 Ω resistor is added in parallel to the inductor to reduce the quality factor.

[7] Two inductor footprints are provided to allow changes. Note that inductor footprints are only present on revision 6 and later motherboards.

[8] This lead spacing is only supported on revision 7 and later motherboards.

[9] Farnell order code for 220 nF: 9751335.

[10] Farnell order code for 22 µH: 1929699.

[11] Starting from revision 7. Motherboards revision 6 and earlier provided only footprints for the Telkor ON-PLC-1 and the PT10 transformer.

[12] For historical reasons, the decoupling capacitor on the secondary side and most parts of the protection circuit are located on the daughterboards. Although footprints for protection components are provided on the motherboard, they are not populated.

Manufacturer	Reference	Designator
Telkor	ON-PLC-1T2-001	TF1
Vacuumschmelze	T60403-K5032-X111	TF3
Wuerth Electronics Midcom	750313480♯	TF4
Oxford Electrical Products	PT10*	TF2

♯ Wuerth can also provide customized versions to meet specific customer requirements.

* Recommended for 6 V operation only.

Table 1: Transformer footprints provided on the motherboard.

2.6 Modem boot sequence

The boot sequence of the NCN49597/9 is controlled by the SEN and TREQ pins (refer to [ON 14a, ON 14b] for details).

For the first, a solder jumper is provided on the daughterboard. By default, daughterboards are populated with the jumper set to ground.

The TREQ (IO2) pin is controlled through the serial port. When the USB cable is connected, TREQ is pulled low.

With these defaults setting the firmware must be loaded into the modem over the serial port.

As an alternative, the SEN jumper can be resoldered to drive SEN high. After each reset, the boot loader in the modem will then load the firmware from an attached SPI memory chip. The memory may must supports the standard commands and must use three-byte addresses. In practice, this will imply a memory size of at least 512 kb as smaller memories typically use two-byte addressing. Refer to [ON 14a, "Boat loader"] for more information.

Starting from revisions c597db-4 and c599db-3, daughter boards include a 2 Mbit LE25U20AQG flash memory.

Alternatively, the user can connect a memory chip to the SPI pins available on the right most daughter board header. Please contact your sales representative for more information.

2.7 Daughterboards

The daughterboard supports modem, transmit and receive filters, protection circuit and the SPI memory. A system-level view is shown in figure 4.

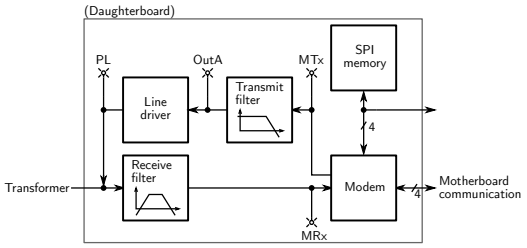

Figure 4: Daughterboard system-level view.

Name	Carrier frequencies [kHz]	Description
D	9–95	CENELEC A-band
E	63.3 & 74.5	IDIS frequencies
G	105 & 115	PL110 frequencies
J	95–150	PL110 2-channel frequencies
K	9–150	CENELEC A–D bands

Table 2: Available daughterboard variants.

2.7.1 Filters

Daughterboards are available for the NCN49597 and NCN49599 with various receive and transmit filter variants. All variants can be used with all firmware stacks.

Each variant is optimised for specic carrier frequencies or an entire frequency band. Table 2 gives an overview (for the filter frequency responses, refer to section 19).

The «E» variant is optimised for the IDIS[13] carrier frequencies (63.3 and 74.5 kHz) widely used for automatic meter reading and similar applications.

[13]IDIS is a specification for smart meters defined by the Interoperable Device Interface Specifications (IDIS) Industry Association.

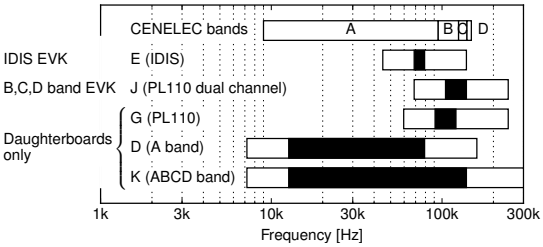

Figure 5: Daughterboards and evaluation kits are available for various frequency ranges. The −1 dB bandwidth of the receive filter is shown in black for each daughterboard, indicating the optimal carrier frequency range. The −10 dB bandwidth is also shown (white).

The «G» variant is optimised for the carrier frequencies of the KNX PL110 standard, 105 and 115 kHz.

The «J» variant covers the B, C and D band, i.e. 95–148.5 kHz.

During field tests it can be very useful to change to carrier frequencies. By choosing frequencies away from interference tones or impedance dips the communication performance can be optimised. For such applications wide-band filters, not optimised for a single carrier frequency pair, are preferable.

«D» daughterboards target the CENELEC A band (9–95 kHz) while «K» cover all CENELEC bands (9–148.5 kHz). The drawback of the larger useable carrier frequency range is the lower suppression of interference.

Figure 5 shows the same information graphically.

«E» are commonly used for automatic meter reading as this band is reserved for utilities in most countries. Other users are typically restricted to the CENELEC B, C and D bands; «G» and «J» filters are then the best choice.

However, it bears repeating that any filter variant can be used provided the correct (i.e. correct for the filter variant) carrier frequencies are selected.

The transmit filter is always a low-pass filter,[14] but the order depends on the variant. Variants requiring three poles can be implemented solely by the two opamps inside the NCS5651 or NCN49599. If five poles are required, an additional opamp (TS972 from ST Microelectronics) is required.[15]

The receive filter of all variants implements a band-pass filter. Typical PLC transmission frequencies are found between 50 and 150 kHz. In this frequency range both active and passive filtering are reasonable solutions.

Revision 3 NCN49599-based daughterboard use passive filtering in the receive path. The topology (figure 6) is fixed; the cut-off frequencies, and therefore the component values, are variant-dependent.

All other daughterboards use active filtering. Again, the topology (figure 7) is fixed with component values changing depending on the variant. The filter is composed of a third-order high-pass filter and a second-order low-pass filter. The filter section built around the buffer amplifier inside the modem realises two poles. The remaining three poles are implemented by an external opamp section.[16]

The frequency response of the available filter vari-

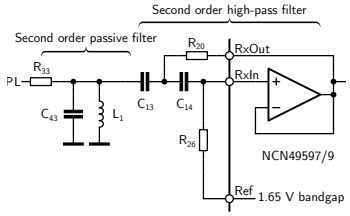

Figure 6: Evaluation daughterboard receive path — passive first filter.

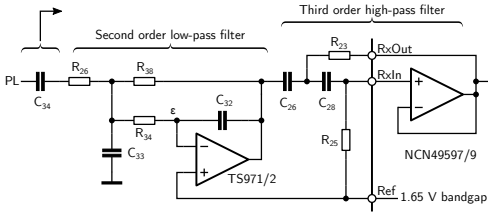

Figure 7: Evaluation daughterboard receive path — active first filter.

[14] Neglecting a single high-pass pole due to the AC coupling capacitor immediately before the first NCS5651 opamp. This capacitor is only intended to shift the DC level from 1.65 V to 3.3 V; it is not intended to filter the signal.

[15] The same printed circuit board design is used for all variants; for the three-pole variants the TS972 is not populated, but is bypassed with 0 Ω resistors.

[16] For variants where the transmit filter requires an additional external opamp, a TS972 is used. For other variants, a TS971 is used. The TS972 is the dual-opamp version of the TS971.

As noted in footnote 15, a single printed circuit is used for all variants. As a consequence, two footprints are required in the receive path. Only one footprint is populated — indeed, only one footprint *can* be populated as they are placed overlapping to save space.

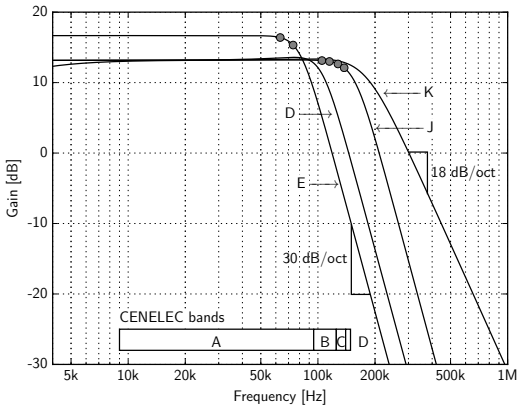

Figure 8: Frequency response of the «D» «E», «J» and «K» daughterboard transmit filters.

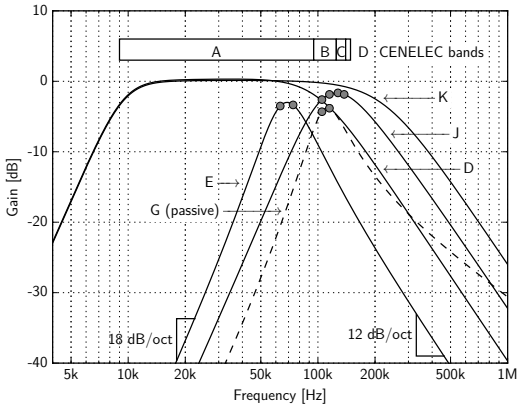

Figure 9: Frequency response of the «D» «E», «J» and «K» daughterboard receive filters.

ants is shown in figures 8 and 9. Only the daughterboard is considered; the gain is calculated from the output of the modem (net `MTx`) to the daughterboard connector pin (net `PL`) resp. from `PL` to the input of the modem (net `MRx`). The coupling circuit on the motherboard is frequency-selective, too; refer to section 2.3 for more information.

In both figures, the IDIS carrier frequencies widely used for automatic meter reading are marked on the «E» trace. The PL110 dual-channel carrier frequencies are marked on the «J» trace.

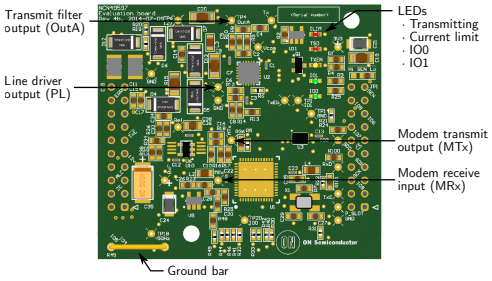

Figure 10: Most important test points on the NCN49597 daughterboard (revision 4).

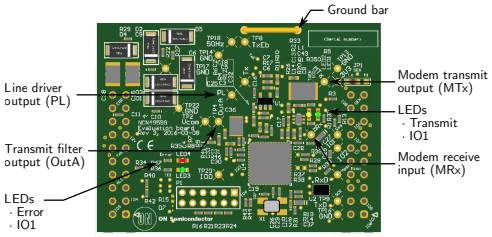

Figure 11: Most important test points on the NCN49599 daughterboard (revision 3).

2.7.2 Using the daughterboards

The various test points on the daughterboard are marked with diagonal crosses. The most important test points are located at the output and input of the modem and at the output of the line driver (refer to figure 4). Called `MTx`, `MRx` resp. `PL`, their location on the NCN49597 (rev. 4) and NCN49599 (rev. 2) daughterboards is shown in figures 10 and 11.

Five LEDs indicate the state of the daughterboard.

The green IO0 and IO1 LEDs are connected to the IO0 and IO1 pins of the modem. The firmware indicates its state on IO0. If the bootloader is running, IO0 is toggled at a slow rate (about 0.5 Hz).

The ON-PL110 firmware toggles IO0 at a faster rate. After it has been loaded the toggle frequency is about 1 Hz. Once configured, the frequency increases — refer to [Ver15b]for details.

The yellow TxEn LED is active when the modem is transmitting.

The red ILIM and TSD LEDs are active when the internal protection circuits of the line driver have detected overcurrent resp. overtemperature. This is normally caused by an excessively low mains impedance.

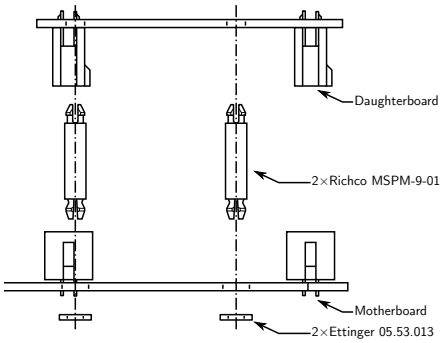

Figure 12: Adding spacers between daughterboard and motherboard improves reliability in environments with heavy vibration.

In contrast to the NCS5651, the line driver of the NCN49599 does not have a thermal flag output; as a result, NCN49599-based daughterboards do not have a thermal shutdown LED.

A ground bar is also provided as a convenient connection point for the ground lead of oscilloscope probes.

2.7.3 Connectors

The two connectors on the daughterboard interface with the motherboard. The left connector carries the «PLC» signals; the right connector the «control» signal to and from the user microcontroller or computer.

The connectors allow convenient swapping of daughterboards and are reliable enough for development.

However, in environments with heavy vibration — for instance close to machinery — the daughterboard can shake loose. To keep it firmly in place, two Richco MSPM-9-01 14.3 mm spacers and two 1 mm washers can be added[17] (figure 12).

Some daughter boards also contain a JTAG header. However, this header is mainly intended for internal development of the modem firmware; it is typical not useful in a user application.

[17] Holes for these spacers are provided on motherboards revision 6b and later; on NCN49597 daughterboards revision 4 and later; on NCN49599 daughterboards revision 3 and later.

Note the «sharp» end of the Richco spacer connects to the daughterboard, while the angled end connects to the washer below the motherboard.

Spacers may be ordered from Farnell (order code 1675869). Ettinger 05.53.013 (Farnell order code 1466903) are suitable washers.

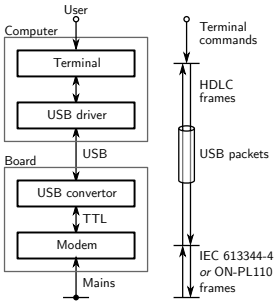

Figure 13: Hardware interfaces (left) and protocols (right) used when controlling an evaluation board from the PLC terminal.

3 Using the evaluation kit

A terminal application has been developed to complement the evaluation boards. The signal flow when using this application is shown in figure 13.

In the terminal, the user types human-readable *terminal commands*, documented below (sections 3.4–3.6). These are translated to binary commands and transmitted to the modem over a serial link. The binary commands are structured according to the widely used HDLC[18] specification and are documented in [Ver15b].

A USB-to-serial conversion chip on the motherboard, working in concert with its computer driver, tunnels this link over USB.

The modem responds to these HDLC appropriately, for instance an incoming HDLC transmission instruction results in an PLC frame being sent. Conversely, upon receiving a PLC frame the modem sends an HDLC notification.

Note the PLC and HDLC frames have a different format; in addition, there is no one-to-one mapping between frame types of both protocols. For instance, a HDLC frame instructing the modem to change the carrier frequencies is not reflected by a PLC frame.

The evaluation kit terminal requires Microsoft Windows XP or later.

One terminal controls only a single board, but multiple boards can be controlled from the same PC by using multiple terminal. To test communication quality over a long distance, two PC's are usually required.

[18] High-level data link control or HDLC is a data link layer protocol standardised as ISO 13239. The standard provices for many use cases, but the modem firmware uses the poin-to-point with asynchronous framing.

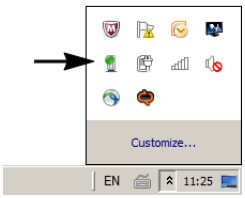

Figure 14: Notification shown by Windows 7 during the automatic installation of the driver of an USB-to-serial converter chip.

Two IEC C14 power cables are required to connect each evaluation board to the mains.

3.1 Driver installation

The serial port is encapsulated over USB as a *virtual communications port* (VCP) by a converter chip. On Windows operating systems, specific drivers must be installed before this converter is recognised by the computer.

On Windows 7 and higher, installation normally starts automatically when the motherboard is first connected. This is indicated by the animated icon in the task bar (figure 14). Note that installation can take a long time the first time an unrecognised device is plugging in — be patient![19]

You can verify the state of the device by opening the *Device Manager* (figure 16).[20] After the driver has been installed successfully and the device has been configured by Windows, a new serial port will appear in the device manager (figure 15) under the category *Ports (COM and LPT)*. The virtual serial port is ready to be used.

If the installation has not (yet) been completed, the device is shown with a yellow icon (as in figure 16, under *Other devices*).

You will need to install the driver manually if the automatic installation fails.

[19] If installation remains slow after the first installation of the driver, consider changing the order Windows uses to search device drivers. This may be done by right-clicking on *My Computer* → *Properties* → *Advanced*→ *Hardware*→ *Device Installation Settings*. Directing Windows to search for drivers on the computer first instead of through Windows Update will speed up device installation.

[20] Accessed by right-clicking on my *[My] Computer*.

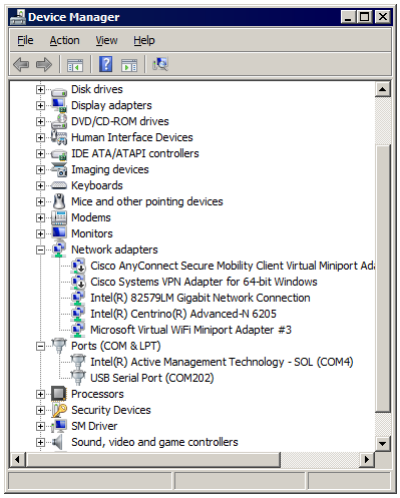

Figure 15: Windows *Device Manager* showing a USB-to-serial converter chip after successful installation.

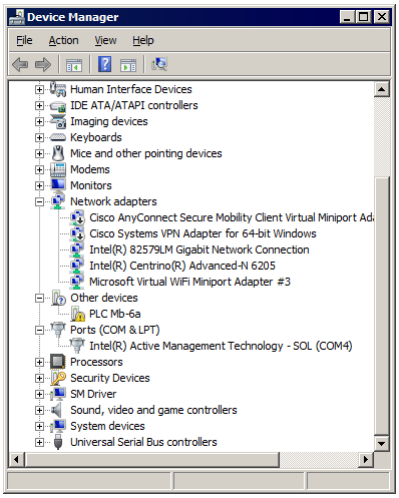

Figure 16: Windows *Device Manager* showing a USB-to-serial converter chip during installation, or after the installation has failed.

- If you have a revision 5 motherboards, with a Microchip MCP2200 converter, you must download the latest version of the driver (MCP2200.inf) from the Microchip website. To install the driver manually, right-click on the device in the *Device Manager*, and select *Update driver software*. Please select «Install from a location on disk» and point the installer to the file you downloaded earlier.

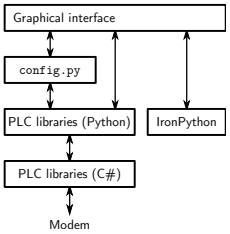

Figure 17: The software architecture of the terminal.

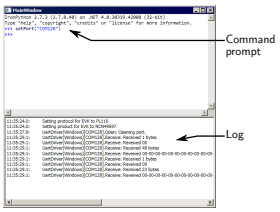

Figure 18: The terminal window.

- All other motherboard revisions use an FTDI converter. The easiest way to install the driver is the installation executable available from [Fut13b]. If you encounter problems during the installation, refer to [Fut09, Fut13a].

3.2 Terminal installation

The terminal is written in C# and IronPython. You will therefore need the .NET framework, version 4 or above. You can download an installer from the Microsoft website at [Mic13].[21] It is strongly recommended to do this even if you have already a .NET framework installed. This ensures your .NET installation is up to date.

You must also install IronPython from `http://ironpython.net/download`.

You should have received an compressed archive with the terminal files. If not please contact `Software.Development@onsemi.com`. Unzip the contents to a new folder — the folder location does not matter.

3.3 Software architecture

Figure 17 shows the architecture of the PC software used to control the modems.

Each PLC modem is addressed over a serial port with HDLC commands. The HDLC commands are generated and parsed by a library written in C#.

To make writing test scripts more convenient, an IronPython application program interface (API) is built on top of the C# library. The API is object-oriented and is ideal for developping complex test scripts involving many modems. Refer to the docstring in the source code for information on how to work with this API; it is not described in this document.

For simple tests, the object-oriented API can be overkill. A second API is made available by the Python module `config`. This API is procedural and is described in subsequent sections.

Both APIs can be accessed from the *terminal* user interface. Start the terminal by executing `PythonTerminal.exe`. A graphical user interface similar to figure 18 should appear. Modems are configured by entering Python statements and expressions in the command entry box. The terminal maintains a history of correctly entered commands. Pressing the up and down arrow keys will cycle through the list.

The user has access to all standard IronPython functionality.

3.4 Getting started with the procedural API

When the terminal is first started, the contents of the Python module `config` are automatically imported. This module contains the procedural API.

The API reference appears in section 5; this section provides an introduction.

When using the procedural API, the terminal application is used to configure *one* single modem and transmit or receive data over the power line.

Connect the modem to the PC with the USB cable.

To find out to which serial port the modem is connected, issue `ports()` in the terminal. The `ports` command lists all available serial ports in the system. Subsequently, unplug the board USB cable and issue `ports()` again. The modem is connected to the serial port that disappeared from the list of available ports.[22]

[21] Due to license restrictions, ON Semiconductor is unable to distribute the .NET installer with the PC software installer.

[22] If nothing changes make sure your operating system recognises the USB-to-serial converter included on the evaluation

Plug the USB cable back in.

To open the serial port, issue the command `openPort` with the port name. For example, use

```
setPort("com213")
```

The port will be opened and associated to the terminal.

To release the port, use `closePort()`[23] command releases the serial port, allowing other applications to use it.

Once the port is properly open and associated to a terminal it is possible to configure the device.

3.4.1 Firmware download

If all previous steps were properly executed the `RxD` LED should start blinking slowly at 0.5 Hz (1 s on, 1 s off), indicating the ROM bootloader is running.

Use

```
downloadLatest()
```

to start the download. This takes a few seconds to complete. Once the download has completed successfully, the `RxD` LED will blink twice as fast (1 Hz). If the download fails, ensure you are addressing the right serial port with the correct baud rate.

3.4.2 MIB parameters

Before the modem can send or receive any data over the power line, it needs to be configured. The modem is configured by settings in the *Management Information Base* (MIB), a set of parameters that control every aspect of the firmware behaviour.

The value of a specific MIB parameter can be retrieved with the `getMibParam` command. To change the value of a parameter, use the `setMibParam` command. The first parameter of both commands is the *parameter name* of the MIB parameter to access. For instance, to change space frequency to 63.3 kHz, use

```
setMibParam("R_FS1", 0x159B)
```

To calculate the MIB value from a desired frequency, refer to section 3.7.

When the space frequency is read back, the same value should be returned:

```
getMibParam("R_FS1")
```

A list of all MIB parameters may be retrieved with the `oids()` command. A detailed description may be found in [Ver15b].

Use these settings to configure your evaluation board as described in section 3.5 (ON-PL110).

3.4.3 Reset

To set all modem parameters back to the default values, use the `reset` command.

Sometimes a hard reset is required, especially if error messages start to pop up in the terminal window. In this case it is recommended to first close the serial port in the terminal with `closePort()`. Then, push the reset button on the evaluation kit itself and restart from from section 3.4.1.

3.5 Using the ON-PL110 firmware

The PL110 protocol is based on a network model without a master node. The protocol supports collision avoidance and detection. Error correction and message reception acknowledgement are also included.

3.5.1 Configuring the modems

Before any data can be sent or received over the power line, the modems need to be configured.

A minimal configuration when working with the MAC layer requires setting the individual MAC and domain addresses and activating the modem. As an example,

```
setMibParam("DLL_ADDR_INDIVIDUAL", 0x20)
```

will set the MAC address to 0020_H .

By default, the ON-PL110 firmware looks for a mains zerocrossing and synchronises transmission with the mains. If synchronisation to the mains zero crossing is not required or not possible (for instance, on a DC network), the MIB parameters `R_MISC_ZCGEN_MODE` *must* be set to 1. If this is not done, the modem will not be able to modulate or demodulate data.

Note that a small performance degradation must be expected (section 3.8).

By default, a mains frequency of 50 Hz is assumed. If 60 Hz is used, the MIB parameter `R_CONF_MAINS_FREQ` must be set accordingly:

```
setMibParam("R_CONF_MAINS_FREQ", 1)
```

board. Refer to section 3.1 for more details.

[23] As a reminder, help is available for all commands are shown in the terminal log window when the command `help(<function>)` is issued.

In addition, the full list of supported commands appears in section 5 on page 19.

Once configuration is complete, switch the modem into active mode:

```
setMibParam("R_CONF_MODE", 1)
```

At least one other node must be configured in order to test power-line communication.

Use a second terminal to connect to a second evaluation board, repeat all steps of section 3.4, and configure the modem with the same parameters, except for the MAC address:

```
setMibParam("DLL_ADDR_INDIVIDUAL", 0x21)
setMibParam("DLL_ADDR_GROUP", 8)
setMibParam("R_CONF_MAINS_FREQ", 1)
```

(for a 60 Hz mains)

```
setMibParam("R_CONF_MODE", 1)
```

3.5.2 Starting basic communication

Once the modems are active, use the command `txmac` in one terminal to transmit communication frames through the data link layer (DLL). The expected parameters are , the destination address, domain address, payload and frame flags. The addresses and flags should be entered in hexadecimal format. As an example,

```
txDll(21, domain=8)
```

will try to send a frame to individual address 0021_H, domain address 8_H with the standard frame format.

By default, each transmitted frame has a payload of 4 bytes; to change this, also specify the `payload` parameter.

It is also convenient to use the broadcast address 0. Broadcast frames are always repeated by the transmitter and the confirmation sent back to the host is always positive. For instance

```
txDll(0, payload=range(10), domain=0)
```

will broadcast a frame of 10 bytes.

If the modem is transmitting, the yellow LED TxEn will light up; if a modem is receiving a correct frame, the green CRC LED will blink. This can help in debugging communication.

The `txDll` function transmits only a single frame.[24] Multiple frames can be transmitted with a Python for-loop:

```
for n in range(10):  txDll(21, domain=8)
```

When another modem receives the transmitted frames and the frame destination address matches the address of the modem, it will report them in the terminal.

[24] Transmitting multiple frames is particularly useful to test the average bit rate, for instance during field tests.

Bit	Name	Value	Description
0–1	Priority	00_B	System
		01_B	Normal
		10_B	Urgent
		11_B	Low
2	Frame type	0	Extended
		1	Standard
3	Repeat	0	Do not repeat
		1	Repeat frame
4	Destination	0	Individual address
		1	Group address

Table 3: Transmission flags of the ON-PL110 `txmac` and `txmacm` commands.

The optional fourth to sixth arguments of the `txmac` function specify the frame options. Together, they form the *frame control flags* byte embedded the PLC frame. The flags, and their possible values, are listed in table 3.

3.5.3 Access to different network layers

It is possible to disable the MAC layer and order the firmware to transmit physical frames directly.

A minimal configuration when working with PHY layer requires disabling the MAC and activating the modem:

```
setMibParam("PHY_DIRECTACCESS", 1)
setMibParam("R_CONF_MODE", 1)
```

Subsequently the command `txphy` can be used; expected parameters are the number of frames to transmit, the payload and bus idle delay. As an example,

```
txPhy((1, 2, 3))
```

will try to send a frame[24] composed of three bytes.

Note that by its nature, reception in PHY mode does not filter frames based on destination address; thus, all frames will be reported by any modem that receives them.

3.5.4 Additional features

The ON-PL110 firmware does not require the zero-crossing detection. This allows it to communicate over DC lines. However, this feature must be specifically enabled with the MIB parameter `MIB_R_MISC_ZCGEN_MODE`:

```
set 4201=1
```

To use dual channel operation, the parameter `MIB_R_CONF_DUAL_CHAN` must be set to 1.

3.6 Advanced terminal features

- The `setLogSeverity(<severity>)` command alters the amount of logging information shown. Use `setLogSeverity("Error")` to disable all output exceps for errors; use `setLogSeverity("Debug")` to shown as much logging as possible.
- The ON-PL110 firmware can report the signal and noise values when frames are received. This can be very helpful during debugging; to enable it, issue `setMibParam("STATS_STREAM_ENABLE", 1)`.

3.7 Changing carrier frequencies

It is evident from the previous section that it is often useful in a field test to change the carrier frequencies.

Ordering the modem to do this is easily done with the MIB parameters `R_FS1` and `R_FM1` (space resp. mark frequencies)[25]. The best-fitting hexadecimal value A corresponding to a given frequency f_C [Hz] is derived from

$$A = ||\frac{4096}{46875} f_C||$$

However, you must also consider the external filters when changing carriers. Both the reception and transmission signal path are filtered. Shifting away from the resonant frequency is liable to decrease the amplitude of a transmitted and received signal. Refer to the designed filter response for your daughterboard variant in section 19 to ascertain the exact result.

Similarly, if a motherboard with a resonant coupling circuit is used (refer to section 2.3), deviating from the resonance frequency will cause signal attenuation[26].

3.8 Attainable communication quality

The practical communication range strongly depends on the network. It is mainly influenced by the attenuation between the nodes and the noise level on the network.

As a rule, high and medium voltage networks have a simple, predictable topology with fairly low noise levels. As electrical loads are not directly connected to the network the impedance seen at the PLC frequencies is high, leading to low attenuation. Therefore, PLC is feasible over long ranges.

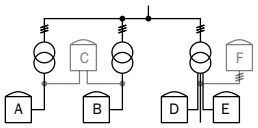

Figure 19: In this example grid, modems A, B, D and E will have difficulty communicating directly.

Conversely, low voltage networks are typically very hostile environments for PLC. This is especially true for residential and office networks. Electrical loads are directly connected to the network. Because many of these loads will have a capacitor across the power lines for EMC compliance, the attenuation at PLC frequencies is high. In addition, the many switching power supplies found in a typical home or office cause significantly higher noise levels.

The complex topology of an office or home network makes it very difficult to accurately predict the communication performance. As a rough guideline, expect up to a kilometre on medium voltage networks and up to a few hundred metres on low voltage networks.

It is important to remember that PLC signals usually cross transformers only with great attenuation.

Thus, communication across phases will often be impossible. In figure 19, modems D and E will not see each other unless the three-phase modem F is configured as a repeater. Note that even in the same room, wall sockets are not necessarily connected to the same phase.

Even when modems are on the same phase, transformers cause problems: in figure 19, modem A and B may not be able to communicate even though they are on the same phase. Adding a two-phase repeater (C) might solve the issue.

In addition, the presence of other consumers can block successful communication. In particular, switch-mode power supplies (SMPS) exhibit a very low impedance at PLC frequencies[27]. In figure 20, the signal sent by modem D to C will be attenuated due to the presence of the television in between.

Referring to the same figure, another aspect be-

[25] The ON-PL110 firmware support dual-channel communication (i.e., four carriers) and therefore also supports the MIB parameters `MIB_R_FS2` and `MIB_R_FM2`.

[26] This is especially true when the mains impedance is low.

[27] The cause of this high load is the presence of a filter capacitor at the input of a switch-mode power supply. The filter is required to comply with conducted EME standards.

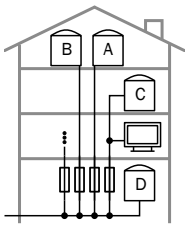

Figure 20: In this sample domestic electrical installation, the television can cause significant deterioration of the communication quality between modems C and D.

comes evident. Even though modem A and B are physically close, communication can be difficult due to the long *electrical* distance. Crosstalk with signals from modem C might also influence the communication.

On an AC mains, a shared time base is available in the form of the mains zero crossing. This time base is used in the PL110 specification to synchronise the modems. The shared modulation clock substantially improves performance.

Thus, when configuring the PL110 firmware for DC operation[28] a drop in performance must be expected. How significant this drop will be strongly depends on the interference characteristics on the mains.

At the same time, in many DC applications the mains is better-controlled than in AC applications. Exact figures depend on the application, but typically, less noise and interference is seen on a DC mains. This balances the performance drop noted above.

4 Troubleshooting

Setting up communication, as described in section 3.5 (ON-PL110 firmware), is a fairly involved process. As a result, an error is easily made.

Section 4.1 describes how to debug *functional* problems; performance problems are discussed in section 4.2.

[28] Note firmware configured for DC operation will not use the zero crossing detection signal, even it is available.

4.1 Troubleshooting functionality

4.1.1 Is the board powered?

The motherboard has two isolated power domains. Both are provided with a green LED indicating the rails of their domain are stable.

Ensure both LEDs are on.

The first LED (located close to the power supply brick) is associated with the daughterboard power domain (3.3 V and 12 V rail). If it does not light up, ensure the mains is connected.

The second LED (close to the USB connector) with the USB-to-serial power domain. If it does not light up, ensure the USB cable is connected and powered.

4.1.2 Is the *motherboard* receiving serial data?

Two green LEDs on the motherboard light up when the USB-to-serial converter is receiving and transmitting data. Send some bytes to the board by issuing a command in the associated terminal (any `getMibParam` command will do, for instance `getMibParam("R_FM1")`). The LED marked PcTx should light up briefly.

If it does not, check the serial communication. Ensure the terminal is associated with the correct COM port.

4.1.3 Has the boot loader started?

After the motherboard has reset (due to power-on, a terminal command, or pressing the reset button on the motherboard), the boot loader should start.

The boot loader indicates it's running by blinking the IO0 LED slowly (1 s on, 1 s off).

If the boot loader is not running, verify with a multimeter the reset and IO2 signal on the daughterboard are both high. Then, reset the modem by pressing the reset button.

4.1.4 Are the modems responding to terminal commands?

Observe the IO0 LED on the daughterboard to see whether the firmware download succeeded.

When the modem first boots, it should blink slowly at a rate of about 0.5 Hz (1 s on, 1 s off). After the ON-PL110 firmware has been loaded, the blink rate increases to about 1 Hz (unconfigured state) or 2 Hz (configured).

As another test, try to read any MIB parameter. For instance, issue the command `getMibParam("R_CONF_MAINS_FREQ")` — the result in the log window should indicate whether the modem is configured for a 50 Hz mains, a 60 Hz mains or a DC network. If the command fails, retry the download procedure described in section 3.4.

4.1.5 Is the transmitting modem actually transmitting?

Once the firmware is up and running, configure one of the evaluation boards as a transmitter[29]. Try to transmit a large number frames and observe the yellow TXEN (transmit enable) LED. During the transmission of each frame, the line driver is enabled and the LED should light up briefly.

If the transmit LED does not light up, the modem is not transmitting frames. The most likely cause is a configuration mistake. Ensure the modem is active (`getMibParam("R_CONF_MODE")` should return 1).

4.1.6 Is the line driver overloaded?

If the transmit LED lights up but communication over the power line is still not possible, observe the red TSD (thermal shutdown) and ILIM (current limitation) LEDs[30]. Any illuminated LED signals the line driver is overloaded.

In contrast to the NCS5651, the line driver of the NCN49599 does not have a thermal flag output; as a result, NCN49599-based daughterboards do not have a thermal shutdown LED. An oscilloscope will be required to check whether thermal shutdown is active.

Monitored the modem output (test point `MTx`) and line driver output (test point `PL`) is useful in any case[31] (figure 21). Short transmission bursts, one for each frame, with a steady voltage in-between are expected (figure 22, top trace). Ensure the sample rate of the oscilloscope is sufficiently high — undersampling will result in gross deviations (bottom trace)[32].

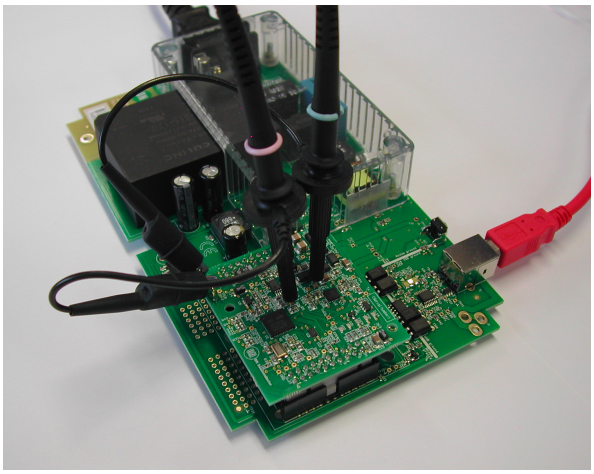

Figure 21: To debug communication problems, observe the `MTx` and `PL` testpoints with an oscilloscope.

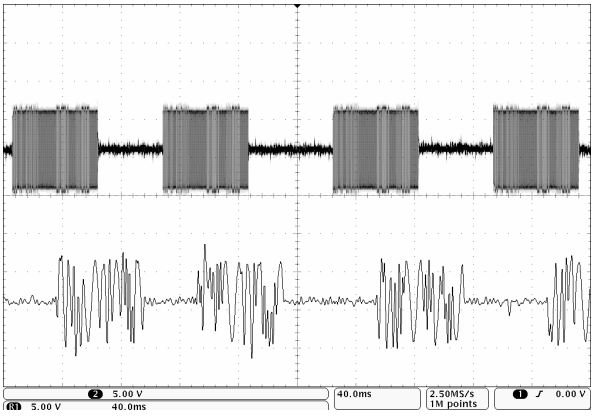

Figure 22: Oscilloscope traces of the line driver output: good (top) and undersampled.

Under normal conditions the line driver output is a close (though amplified) copy of the modem output (figure 23[33]). Note the amplitude at the line driver output changes slightly between mark and space — this is caused by slightly different gains in the transmit

[29] To do this, refer to section 3.5.1.

[30] Note that the NCN49599 modem does not have a thermal shutdown indication pin, and no thermal shutdown LED is provided on NCN49599 daughterboards.

[31] The signal `MTx` is a small-signal, low-power signal generated by the 3.3 V analogue part of the modem. Therefore the amplitude is small (1 V_P with a 1.65 V DC bias). This signal is not output directly on the mains, but is amplified by the line driver to the signal at `PL`. The DC bias here should be 6 V with an amplitude of approximately 10 V_{PP} (neglecting any configured transmit attenuation). Refer to figure 4 for an overview.

[32] Both traces were not captured at the same time, resulting in a time offset.

[33] The figure shows an oscilloscope screen shot, split in two parts. The top part shows the entire trace length: two traces are shows (appearing black and grey). The bottom part shows the same traces but zoomed-in. This oscilloscope feature is called «Wave Inspector» by Tektronix and «MegaZoom» by Agilent.

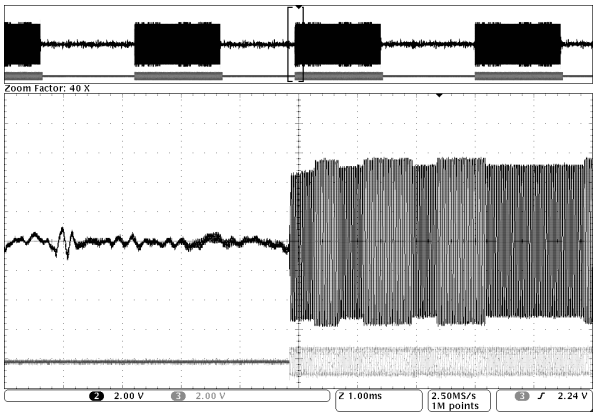

Figure 23: Modem output (bottom trace) and line driver output (top trace) under normal conditions.

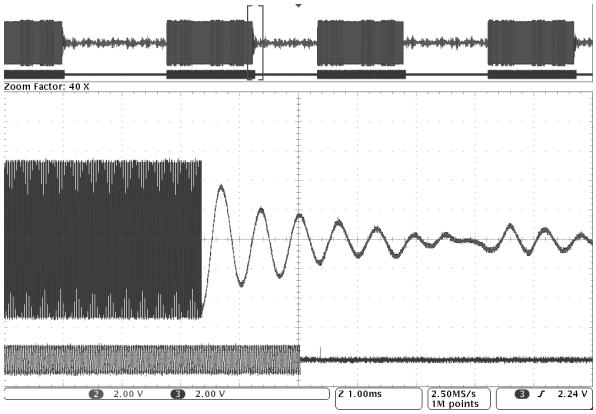

Figure 24: Modem output (bottom trace) and line driver output (top trace) during thermal overload.

filter and is perfectly normal.

If the signal at the line driver output suddenly disappears before the frame is finished, as shown in figure 24, thermal shutdown is a likely cause.

If the output is heavily distorted, the current limitation may have been triggered. This can happen if the line impedance is too low.

To fix the overload, reduce the transmitted amplitude with the `R_ALC_CTRL_VAL` MIB parameter:

`setMibParam("R_AGC_CTL_VAL", 1)` (or any value between 0–7)

This parameter sets the number of 3 dB attenuation steps, 0 corresponding to 0 dB attenuation, 7 to 21 dB attenuation. Try to maximise the transmitted amplitude, without overloading the line driver.

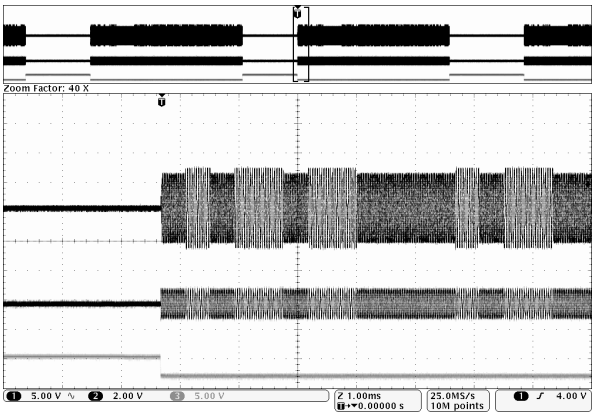

Figure 25: Modem output (centre trace) and a clipped line driver output (top trace). The line driver enable signal is also shown (bottom).

If a well-formed signal is seen on `MTx` but clipping is observed on `PL`, the usual cause is the choice of carrier frequencies. Carrier frequencies must be chosen appropriately for the transmit filter of the transmitting daughterboard. Refer to table 2 and section 19 for a list of daughterboard filters. If required, change frequencies (section 3.7).

Once this is done, the transmitter is set up. The next step is to check the receiver.

4.1.7 Is the transmitted signal clipped?

If the gain settings are not appropriate for the chosen carrier frequencies, the output signal can be clipped. This will often (but not always) overload the line driver. Communication is still possible with moderate clipping, but for a more robust design it should be avoided.

Clipping is not evident with large oscilloscope time bases. For instance, the signal at the transformer secondary in figure 25 (top trace) looks fine at 1 ms/division, but when the time base is decreased to 20 µs/division, (figure 26) it transpires the gain is too high at the mark carrier frequency.

4.1.8 Is the receiver correctly configured?

When the transmitter is functional, but the terminal associated to the receiving evaluation kit does not report that frames are being received, verify the receiving modem is configured correctly.

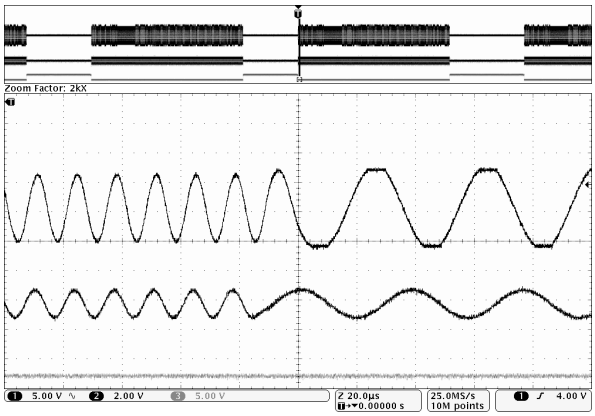

Figure 26: Modem output (centre trace) and a clipped line driver output (top trace). The line driver enable signal is also shown (bottom). Note the time base.

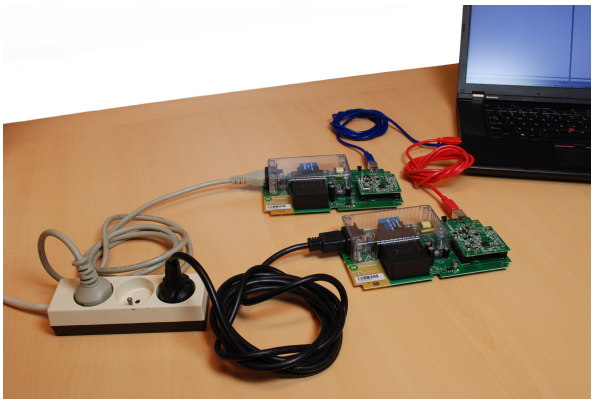

Figure 27: Connecting both modems to single distribution strip reduces the noise levels and attenuation.

Of course, the same carrier frequencies must be used by all modems.

When using the ON-PL110 firmware, the media access control (MAC) layer address and domain must be set. They must also match the destination specified in the `txDll` command sent to the transmitter.

4.1.9 Does the receiving modem see a strong signal?

Connect the two evaluation boards to a single power distribution strip (figure 27) — this lowers the attenuation from transmitter to receiver. Also, do not connect other consumers to the distribution strip — this lowers the local noise and interference level.

When it has been verified the transmitted is generating a strong signal the PL test point (section 4.1.6), observe the same test point on the receiver with another oscilloscope probe. An identical, somewhat attenuated, waveform should be seen.

If no signal is seen, the most likely cause is excessive attenuation in the coupling circuit[34]. Ensure appropriate carrier frequencies are selected for the type of motherboards[35]. Resonant motherboard only operate well with carrier frequencies close to the resonance frequency. For more details, refer to section 3.7.

Once a good waveform is seen on the test point PL on the receiving daughterboard, observe the signal at the input of the receiving modem (test point MRx). Here, too, a fairly strong signal is expected (about 1 V_{PP}). If a strong signal is seen on PL but not on MRx, the carrier frequencies are not chosen correctly for the receive filter of the receiving daughterboard[36].

With an acceptable signal on the receiver test point MRx, communication should be possible.

If nothing helps, consider reversing the role of receiver and transmitter.

4.2 Troubleshooting performance

Sometimes the communication over the power line between two evaluation kits is functional but low-quality, i.e. the receiving modem receives some but not all frames.

As in section 4.1.9, connect all modems to a single power distribution strip. This should ensure most frames are received correctly. If not, observe the test points PL (on the motherboard; transformer secondary) MRx (on the daughterboard; modem input). Severe noise and interference at this point can interfere with the communication carriers.

When the communication quality is disappointing in the field, it is very useful to study the local topology of the grid to understand the underlying causes. Some

[34] Because the modems are connected on the same power distribution strip, the attenuation for PLC frequencies between the power outlets should be low, even in fairly heavily loaded environments.

[35] Also ensure the receiving modem board is not transmitting, thus blocking the signal from the other board. The state of the board (transmitting or receiving) is indicated by the TXEN LED.

[36] Again, refer to section 19 for a list of daughterboard filters.

clues where to look to try to understand the behaviour can be derived from section 3.8:

- Are the evaluation kits on the same phase? If not, is it possible to add a repeater connected to both phases?
- How does the electrical installation plan look? Frequently, such a plan will be difficult to find. If the cabling between two nodes is long, is adding a repeater feasible?
- How many consumers are connected to the network, in particular those with heavy switching power supplies located between the evaluation kits? Does the communication quality improve when disconnecting consumers?
- When two modems that are close-by (i.e., seeing little attenuation between them) are both configured as repeater, slight timing difference will cause the repeated frames to interfere destructively. Thus, the addition of a second repeater function will actually cause *worse* performance. In such a case, try configuring only one of these modems as a repeater.
- What disturbance is seen on the mains? In particular, interference on or near the carrier frequencies will significantly reduce the communication quality. However, interference away from the carrier will also cause problems if the amplitude — after filtering — is sufficiently high.
  Checking the waveform and spectrum of the signal at the modem input (test point `MRx`) will help to pinpoint this[37].
- What is the attenuation for each carrier frequency? Are the carriers by chance located at a notch of the signal transfer function of the network?

5 Procedural API reference

The following commands are related to the terminal itself and its serial port. They are always available.

`setLogSeverity(severity)` Change the log level.
The current log level determines the verbosity of the log output. All log messages with a level higher than or equal to the current level are printed.

- `severity`. New low level; should be one of "Error", "Warning", "Info", "Debug".Change the logging level. All messages with the given severity or above will be logged.

`setPort(port, baudrate=115200)` Opens a serial port on the PC with a given baud rate.

- `port`. The name of the COM port (e.g. COM13)
- `baudrate`. The baud rate, depends on the daughterboard. Refer to page 12 for details.

`closePort()` Closes the port that is currently open.

`downloadLatest()` Download the latest PL110 firmware to the modem.
Send the latest ON-PL110 firmware as a binary image over the UART into the modem. Prior to using this command, the serial port must be open. This command will control the TREQ pin for you; the SEN pin must be manually driven low for the download to succeed. Refer to section 2.6 on page 6 for details.

`setMibParam(mibParam, value)` Retrieve the value of a given MIB parameter.
Read the specified Management Information Base (MIB) parameter from the modem. The result is returned as an integer. For instance, `getMibParam("R_FS1")` might return 5531 (corresponding to 63.3 kHz).
Note that some MIB parameter are not writeable if the modem is configured.[38]

- `mibParam`. The MIB parameter to retrieve; should be given as a string, e.g. `"R_FM1"`.

`setMibParam(mibParam, value)` Change the value of an MIB parameter.
Change the value of a given Management Information Base (MIB) parameter of the modem

[37] A spectrum analyser may be used for this, but a digital oscillope with (offline or online) FFT is often more convenient. Note that an oscilloscope is severely limited in dynamic range compared to a spectrum analyser.

[38] The modem is configured once the MIB parameter `R_CONF_MODE` (4102_H) has been set to 1. This can be done only once; reset the modem and re-upload the firmware to revert to an unconfigured state.

to the specified value.
For example, `setMibParam("R_FS1", 0x159B)` changes the space frequency to 63.3 kHz.

- `mibParam`. The MIB parameter to change; should be given as a string.
- `value`. The new value for the parameter as an integer.

`txPhy(payload=range(30), frame_count=1)`
Transmit a frame at the Physical Layer level.
Transmit a single frame with the given payload. The payload should be specified as a sequence of integers. The modem should have been configured to operate at the physical layer using `setPhy()` or `setMibParam("PHY_DIRECTACCESS", 1)`.

- `payload`. Data to transmit, sequence of integers.
- `frame_count`. Number of frames to transmit.

`txDll(destinationAddress, payload="ABCD", domain=8, repeat=True, priority="low", ext_frame=False,`
Transmit a frame at the Data Link layer (DLL) level.
Transmits a MAC frame from the modem by consecutively sending a number of Data Requests to the modem. The requests contain the data requested by the user. Whether or not the transmission was successful is not regarded by this command; that is left to the user.

- `destinationAddress`. Individual address of the destination node, an integer.
- `payload`. Data to transmit, sequence of integers or a string.
The first 8 bits control the handling of the frame by the MAC layer; the format is dependent on the frame type. For standard frames, the lower three bits (bits 0–2) specify the hop count. For extended frames, bits 0–3 specify the «extended frame format» (EFF); bits 4–6 specify the hop count. All other bits are reserved and should be zero.
The bytes specify the user data. In KNX applications the first byte of the user data block is the «transport layer protocol control information» (TPCI) field.[39] For user-defined applications the first byte can be used as any other data byte.
- `domain`. Domain of the destination node, an integer.
- `repeat`. Repeat this frame.
- `priority`. Frame priority, one of "low", "normal", "urgent" or "system".
- `ext_frame`. Send an extended frame.
- `frame_count`. Number of frames to transmit.

`setDll(localAddress)` Configure the modem to operate at the Data Link Layer (DLL) level.
This is a convenience function setting a minimal configuration.
The individual address is configurable; the domain is fixed at 8.

- `localAddress`. Address of the modem in the PL110 network, an integer.

`setDllDc(localAddress)` Configure the modem to operate at the Data Link Layer (DLL) level in DC mode.
This is a convenience function setting a minimal configuration.
The individual address is configurable; the domain is fixed at 8.

- `localAddress`. Address of the modem in the PL110 network, an integer.

`setPhy()` Configure the modem to operate at the physical Layer («phy») level.
This is a convenience function setting a minimal configuration.

5.1 Useful Python commands

All functionality of IronPython is available in the terminal; refer to the IronPython documentation for a complete reference. Some useful functions when writing scripts are

[39] To quote the KNX standard, «The TPCI controls the Transport Layer communication relationships, e.g. to build up and maintain a point-to-point connection.».

`print content [, content]` Outputs text to the terminal window.

- `content`. Data to print.

`time.delay([seconds])` Pause the script execution for some time.

- `seconds`. The pause time in seconds (can be fractional).

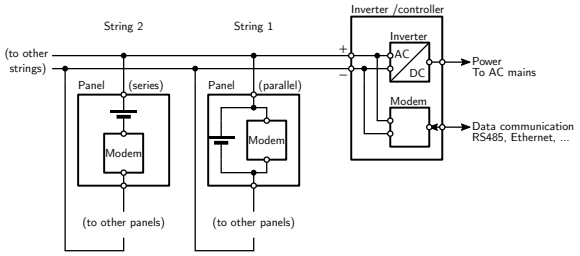

Figure 28: Typical solar panel topology.

Part II
Application design manual

The evaluation kit is not a reference design: the designer must adapt the schematics and layouts to fit the requirements of the end user. Your ON Semiconductor sales representative can help you with this. This section gives a few hints.

6 Theory of operation

6.1 Communication speed

The raw communication speed is quoted in *baud*, i.e. symbols per second. With S-FSK, one symbol represents one bit, therefore the raw baud rate is identical to the raw bit rate (bits per second).

However, only a part of this capacity is available to the user, for three reasons.

The user data is packaged in *frames*, adding information required for demodulation (the preamble), routing (destination addresses) and the protocol (checksum &c.). The overhead introduced by this depends on the protocol.[40]

Additional capacity losses stem from higher network layers. In the ON-PL110 protocol, nodes can only transmit a number of bit times after the last detected frame has finished. Acknowledgments also take up network capacity.

The third reason for capacity reduction comes from collisions. Because ON-PL110 nodes can start transmitting simultaneously, frames may be lost. The impact of this is highly complex and depends on the network congestion. A network close to 100% capacity (taking the above into account) will see many collisions, leading to low throughput of user data despite the high bus activity.

On a 50 Hz or DC mains, ON Semiconductor modems can communicate at 300, 600, 1200, 2400 or 4800 baud. 9600 baud is possible with «dual channel» operation, but this is still in an experimental stage.

On a 60 Hz mains, the baud rates quoted above should be multiplied with 1.2.

6.2 Signal propagation — traditional

6.3 Signal propagation — solar panels

The topology of a typical solar panel is both simpler and more complex than the traditional AC mains topology.

It is simpler because the topology is predictable. The number of possible configurations is limited, especially if the inverter and the solar panels are made by the same manufacturer.

The signal injection is however more complex. A typical solar panel application has multiple string of series-connected panels (figure 28). Each string is connected in parallel to the inverter inputs.

To couple the modem inside the solar panel module to the mains, two solutions can be discerned.

The first is to place the coupling in series with the panel electronics (or panel alone in case of a passive panel). This is called «series coupling». In this case, care needs to be taken that the PLC signal can pass through the panel electronics.

The second solution places the panel electronics in parallel to the modem («parallel coupling»). The panel electronics must now be designed to ensure that the PLC signal *cannot* pass through, as this would bypass the modem.

[40] As an example, the ON-PL110 protocol adds 9 bytes to each standard frame [Ver15b, tables 4.3, 4.4 and 4.5].

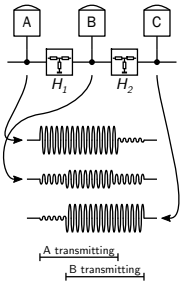

Figure 29: Simultanenous transmissions from modems may interfere, depending on the mains environment. Signals are shown for $H_1 \approx H_2$.

6.4 Co-existence

PLC systems communicate over a shared medium. As with any channel-sharing network, there is a fundamental risk of a transmission of one node interfering with another.

A modem will be unable to receive a transmission from another modem any time the signal of this transmission *at the receiving modem* is swamped by a third-party transmission. The relative signal strength at which demodulation becomes impossible is difficult to predict, as it depends on the timing (when the interference starts and stops), the general noise level, and the absolute signal levels. A reasonable rule of thumb is that demodulation starts failing when the level of the desired signal is less than 10–20 dB stronger than the interfering transmission.

Figure 29 shows a topology where this can occur. Assume modem A is transmitting a frame to modem B, and modem C starts transmitting as well. Which signal is dominant at modem B depends on the mains environment, in particular on the attenuations H_1 and H_2. If $H_1 \approx H_2$, modem B will not be able to receive either transmission. If $H_2 \gg H_1$, modem B will correctly receive the transmission from A if it started first,[41] but not from C. For $H_1 \gg H_2$ the opposite holds.

As an obvious consequence, a modem can not receive while it is transmitting; in real-life networks the transmitted signal will be at least a strong as any received signal.[42]

Because of lack of full-duplex communication at the physical layer, higher layers must implement collision avoidance. The exact mechanism depends on the communication standard.

In the IEC 61334 protocol, inside a network a strict master-slave topology is used. The master is in full control when an individual slave will transmit.[43] However, collisions across networks (each network having a single master) can still occur. In practice this is minimised by the careful partioning of the network, ensuring a high attenuation between networks.

The ON-PL110 protocol uses mesh networking. Nodes do not need permission from a master to transmit. The protocol specifies *collision detection*, allowing a transmitting node to . The transmission can be repeated at a later time.

Refer to [Ver15b, Ver15a] for detailed information on these mechanisms.

7 Coupling circuit

The coupling stage is the physical interface between a PLC modem and the mains. Its impact on the overall performance is substantial.

The coupling circuit varies from design to design since the requirements in different markets differ significantly.

The designer must first chose between *inductive coupling* and *capacitive coupling*.

Inductive coupling inserts a current transformer in series with one of the mains lines[44] (figure 30). The transformer can not superimpose the PLC signal between line and neutral; it can only try to force a current in one line. The network topology must be amenable to this. Typically, communication is only practical between nodes placed in a string on a single line.

Because of these limitations, inductive coupling is infrequently used outside speciality systems such as

[41] If the transmission from C started first, the results are unpredictable and will depend on the absolute signal levels.

[42] Note this is a general property of power line communication, not a limitation in the PLC modems from ON Semiconductor. However, *because* it is a general property, the reception and transmission paths of the modems have been designed to share some blocks in order to reduce cost. The modems *in se* are therefore by design not full-duplex.

[43] The Linky specification adds a measure of full-duplex communication in the form of *alarms* [ERD09]. Through an alarm, a slave can communicate uncommanded to a master. Alarms are however a minor part of the communicated data and do not change the essential master-slave characteristic of IEC 61334.

[44] Not to be confused with the transformer-isolated coupling discussed below; despite the impedance transformation and the isolation provided by the transformer, the actual mechanism to couple the PLC signal onto the mains is still capacitive.

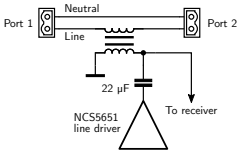

Figure 30: Basic outline of an inductive coupling circuit.

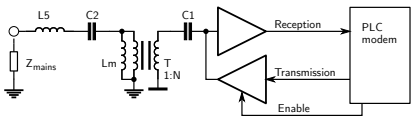

Figure 31: Basic outline of a transformer-isolated coupling circuit. L_M represents the magnetisation inductance of transformer T.

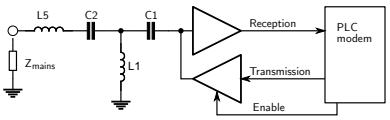

Figure 32: Basic outline of a non-isolated coupling circuit.

medium-voltage lines and solar panel control.

The vast majority of applications use capacitive coupling stages. These can again be divided into two groups: transformer-isolated (figure 31) and non-isolated (figure 32) designs[45].

In both figures, the components have a similar function. We will discuss them for the most common case of coupling to an AC power line.

C2 is a high-voltage (X2 rated) capacitor, which must pass the PLC signals (carriers) while blocking the high voltage fundamental (DC/50/60 Hz) from the mains. Some 50/60 Hz current is inevitably injected through C2.

To lower the impedance of C2 at PLC frequencies, a series inductor (L5) is sometimes added, forming a resonant tank with C2. This comes at the cost of a lower bandwidth, constraining the choice of carriers to frequencies close to the resonant frequency.

C1 is used to block the DC voltage of the output driver[46]. As with C2, it must not generate noticeable distortion at PLC frequencies.

An inductive component in parallel with the signal chain is needed to shunt the 50/60 Hz injected current. At PLC frequencies its inductor impedance must be high.

Similarly, an additional capacitor can be added in parallel to L1 or T to create a resonance at PLC frequencies. Again, the coupling impedance is decreased (as the impedance of L1/T[47] is compensated) at the cost of a lower bandwidth. Because it makes for a very difficult design procedure, ON Semiconductor recommends against adding such a capacitor unless the effects of tolerances are exhaustively studied[48].

Note that the circuit behaviour changes when switching from reception to transmission mode. In a typical design, the impedance of the net to the right of C1 drops to a fraction of an ohm during transmission when the transmit amplifier is enabled.

7.1 Transformer-isolated coupling

A transformer-isolated coupling (figure 31) is used if mains impedance matching or isolation between modem and mains is required.

Impedance matching is useful (or required) to drive very low mains impedances. The power line impedance varies consirably with region, location in the distribution network, time of the day and frequency.

Commonly observed impedances of $1 \ldots 5\ \Omega$ would require the driver to deliver a high current at a low voltage. Usually, increasing the signal voltage and decreasing the current results in a more efficient system. The transformer increases the mains impedance seen by the driver[49].

A list of suitable transformers appears in table 1. The application note [VC11] on transformer design provides further information.

Note impedance matching is not intended to reduce reflections as in RF circuits. At typical S-FSK

[45] Although vital to the survivability of the circuit, protection circuits are not shown in there figure for clarity. For more information, refer to section 7.4.

[46] In the — uncommon — topology with dual supplies or a push-pull driver this capacitor may be omitted provided the output voltage offset of the line driver is low enough. A substantial offset can cause core saturation in the inductive shunt element (T/L1).

[47] More specifically, the impedance of the magnetisation inductance L_M.

[48] Note however that added a small capacitor (order of 1 nF) to reduce the effects of current spikes is strongly recommended. Refer to section 7.4 for details.

[49] It is important to keep in mind that since transformation applies to both current (1:N) and voltage (N:1) the impedance will be scaled N^2 times — 4 times with a 2:1 transformer (driver current will be doubled while driver voltage will be halved).

frequencies, reflections are not a problem. In another contrast to RF design practice, it is rarely desired to match the coupling circuit impedance to the mains impedance. Although this would maximise the power transmitted, it is not efficiency-optimal. Instead, the design goal usually is to minimise the coupling impedance to maximise the amplitude on the mains.

7.2 Component selection

To design the coupling circuit, the two most important parameters are the carrier frequencies and the tolerable mains impedances at the carrier frequencies.

Other requirements also apply. Many standards require a minimum impedance of the device as seen from the mains[50].

That said, the main driver of the design is keeping the impedance of the coupling stage low during transmission.

Selecting a component for C1 is usually easy. Because the voltage rating is low (the supply voltage) high capacitance values (in the order of 10 µF) are readily available without burdensome cost. This achieves the main design objectives of negligible distortion and low impedance.

The magnetisation inductance of T[51] (L_M) must block the PLC signal and shunt the 50/60 Hz voltage. An inductance in the order of 1 mH is appropiate, yielding an impedance of 310 mΩ resp. 310 Ω at 50 Hz resp. 50 kHz.

To lower the impedance of the coupling circuit, it is desirable to select a high value for C2. However, this will result in a higher residual 50/60 Hz current through L1/T. It is important that the magnetics do not saturate with this current and the PLC signal combined, i.e.

$$I_{\text{sat},L} > \frac{\sqrt{2}V_{\text{mains,rms}}}{2\pi \cdot C_2 \cdot 50\,\text{Hz}} + I_{\text{PLC},Pk} \qquad (1)$$

with $I_{\text{sat},L}$ the current through T still not resulting in unacceptable distortion[52] and $I_{\text{PLC},Pk}$ the peak current of the PLC signal. Note that the latter is dependent on the mains impedance and the coupling circuit.

From the lowest mains impedance to be supported and an estimate of the coupling impedance, $I_{\text{PLC},Pk}$ may be estimated. From equation 1 the maximal value of C2 can then be derived.

If a resonant circuit is desired to lower the impedance, the value of L5 can be calculated using

$$L_5 = \frac{1}{C_2 \left(2\pi \cdot f\right)^2}$$

with f the resonant frequency; it is best chosen between the carrier frequencies. The saturation current of L5 must be chosen according to equation 1 to avoid distortion.

It is clear the design flow is not trivial; usually a few iterations are required to achieve an acceptable solution.

7.3 DC coupling

So far, we have discussed AC applications. If a coupling circuit to a DC line must be designed the same principles apply. In fact, a coupling circuit designed for an AC application will work perfectly under the same circumstances[53] on a DC line.

However, a custom coupling stage can be less expensive.

First, decide on the minimal impedance at PLC frequencies to be supported. This will allow deciding between an transformer-isolated and a non-isolated coupling (unless the isolation itself is a requirement).

The value of L1/L_M should be chosen on the same principle as before: its impedance should be high at PLC frequencies, and low at the expected frequencies of dominant interference.

Saturation of L1/T due to currents resulting from interference is usually not a concern. Therefore the core of these components can often made smaller (and cheaper).

In a non-isolating coupling stage, C1 and L1 may be omitted althogether provided the amplitude·frequency product of the main interference is low enough. This results in the circuit depicted in figure 33.

If the maximum DC voltage over C2 is low, high capacitance values for C2 are not expensive. To achieve

[50] A typical example is the widely adopted standard EN 50065; refer to [EN01b] for details.

[51] In the following, we will discuss the isolated case. However, the same applies for the non-isolated case, provided L1 is substituted for L_M.

In addition, calculation will be done for 50 Hz only.

[52] Note that magnetics manufacturers define saturation current as the current resulting in a 10% or 5% inductance drop. For PLC applications, as such a large drop will result in excessive distortion. Therefore $I_{\text{sat},L}$ is usually much smaller than the manufacturer-specified value.

[53] These circumstances include carrier frequencies, line impedance and interference from the line.

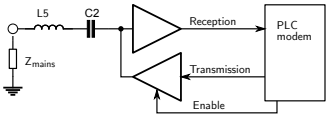

Figure 33: Basic outline of a simple non-isolated coupling circuit for DC lines.

a low impedance for this circuit, one can rely again on the resonance of L5 and C2. However, it can be cheaper to increase the value of C2 (as an example, 20 µF $\sim$ 0.15 Ω at 50 kHz). The distortion added by the capacitor must be taken into account.

Note that such a circuit is very susceptible to transients and spikes from the DC line. Retaining L5, with a smaller value, is strongly recommended for protection; adding a ferrite bead may be a good alternative. Even then, this solution is likely to be cheaper than a resonant inductor.

If a very low output impedance is not required, L5 may be replaced with a resistor in the order of 2 Ω.

7.4 Transients protection

High-energy transients and spikes occur commonly on the mains network. Any circuit coupled to the mains must be protected from these. A PLC modem is particularly vulnerable because a high-frequency path to the mains is required for communication.

In addition, the system must also withstand hot insertion[54]. Hot insertion is particularly troublesome because the DC supply rail is still ramping up. For instance, a positive spike from the mains will be clamped be clamped by a 6.8 V zener diode above 6.8 V only — but when the supply rail of the modem has only ramped up to 3 V at the moment the spike occurs, an IC will still see a voltage differential of 3.8 V over its internal ESD protection diode. Therefore, the protection scheme must not be limited to fixed-voltage components such as a TVS and a zener diode.

The expected transients on the mains will drive the design of the protection circuit.

The recommended protection circuit for connecting to a low-voltage grid is shown in figure 34. Voltage spikes from the mains are successively absorbed by a metal oxide varistor (MOV), a transient voltage suppressor (TVS) and Schottky diodes.

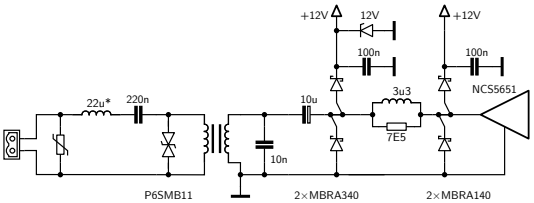

Figure 34: Recommended protection circuit for the NCS5650.

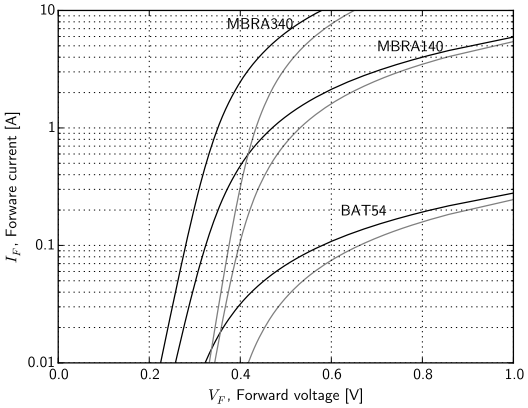

Figure 35: Forward current versus forward current for (top to bottom) MBRA340, MBRA140 and BAT54. Curves are shown for 25°C (black) and −40°C (grey).

The circuit make good use of the low voltage drop (V_D) of Schottky diodes — typically 0.4 V, about half of the 0.7 V voltage drop of silicon diodes. However, care must be taken to select devices with a suitable rated current. During transients, a very short but high current flows through these diodes; the voltage drop must remain low even for those elevated currents. Small-signal diodes such as the BAT54 are wholly unsuitable (as illustrated in figure 35); instead, diodes such as MBRA140/340 (rated at 1 resp. 3 A) must be used for adequate protection[55].

The current flowing through the Schottky diodes to ground in figure 34 is not a cause for concern provided the layout is done thoughtfully (see below). However, current flowing through the diodes to the 12 V supply

[54] Although an automatic metering system is usually always-on, the meter still must survive the initial installation and power outages.

[55] Even larger diodes do not bring much additional benefit, but their cost is much higher.

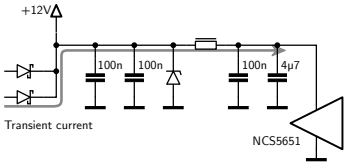

Figure 36: Recommended connection to the power supply pins of the line driver.

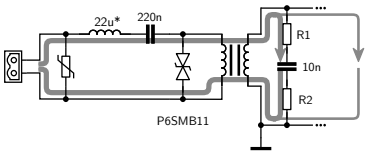

Figure 37: Parasitic resistance at the 10 nF capacitor in figure 34, and the relevant spike current flows: desired (thick grey) and unwanted (thin grey).

rail will increase the voltage on this rail somewhat. Unchecked, this may lead to an overvoltage on the supply pins of the line driver.

Figure 36 shows the recommended solution. A zener diode and capacitors[56] are placed close by the diodes to the supply line. They divert the spike current to the ground, allowing it to flow back to the transformer winding.

To protect the amplifier from the remaining voltage rise, a ferrite bead is placed between the supply line and the amplifier supply pins. The ferrite bead must be rated for the maximal current expected during transmission. Ceramic capacitors must be used at the amplifier pins.

When adapting the recommended protection circuit for an application, the tolerance of the power supply voltage must be taken into account.

The designer must also consider the temperature range of the circuit; the temperature coefficient of the TVS and the zener diode must be taken into account[57]. Similarly, V_D increases as the temperature drops.

The layout of the protection circuit is critical due the high-frequency content of spikes from the mains. An uninterrupted ground plane is essential in this area. Although the protection circuit is usually drawn similar to figure 34, figure 38 might be useful as a layout-oriented schematic. It is important to design the layout for minimal impedance in the loops the designer wants the spike currents to follow. Consider for instance the 10 nF capacitor directly behind the transformer. In figure 37 the parasitic resistances at the terminal of this capacitor and flow of the spike currents are shown[58]. The desired path (thick grey line) flows through the capacitor. Inevitably, some current flows through the remainder of the circuit (thin grey). By minimising the parasitic resistances, this current can be reduced.

Figure 39 shows two possible layout of the 10 nF capacitor; at the left, the connection from the ground and signal tracks are kept very short; indeed, the track runs through the pad of the capacitor. This reduces the impedances R1 and R2 of figure 37 to the absolute minimum.

On the other hand, the layout on the right side introduces a stub track to the signal track. This causes parasitic resistance and inductance, which in turn causes the spike current to flow increasingly in the remainder of the circuit. The connection to the ground signal is even worse: instead of first routing the ground track from the transformer to the capacitor and only then to the ground plane, the capacitor pad is connected with a via to the ground plane. This via adds significant inductance (included in R2 in figure 37).

All other current loops should be studied similarly, as the same considerations apply. It must be noted in particular that the diodes to the 12 V supply must be provided with a ceramic decoupling capacitor close by. The spike current flowing into these diodes can *not* flow into the normal decoupling capacitor because these loops are typically too large, and offer too great an inductance to this current; a current that is the epitome of a high frequency signal indeed.

[56] The zener diode and the capacitors are also shown in figure 34.

[57] As the temperature falls, so does the zener voltage. Note that this does not hold for parts with a low zener voltage ($\lesssim$4 V), for which the temperature coefficient is negative.

As an example, we will assume an operating range of $-35°C$ to 90°C; a power supply of $12\,V \pm 10\%$; and a zener voltage coefficient of $5\,mV/°C$. In this case, the minimal nominal zener voltage should be $12\,V \cdot 110\% - (-35°C - 25°C) \cdot 5\,mV/°C = 13.5\,V$. A part with the next higher available voltage should then be selected. Note that the supply voltage has no influence on the TVS nominal voltage: this voltage is determined by the transmitted signal amplitude.

[58] Current also flows through the MOV and the TVS; these are neglected here to simplify the example.

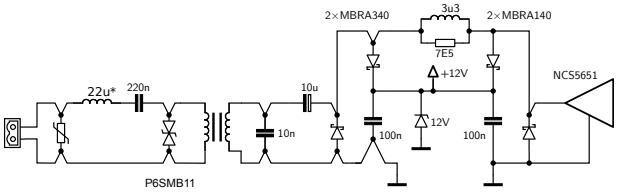

Figure 38: The circuit of figure 34, redrawn to clarify the recommended layout.

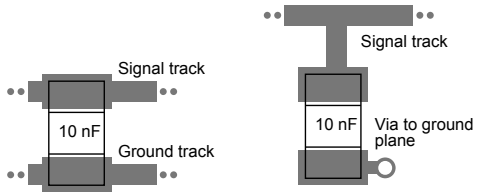

Figure 39: Possible layouts of the 10 nF capacitor in figure 34: good (left) and bad.

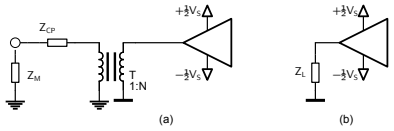

Figure 40: Simplified model of the coupling circuit of figure 31.

8 Line driver operating area

Consider the typical line driver and coupling circuit depicted in figure 31. Obviously, the PLC signal amplitude and power that can be superimposed on the mains is limited. In a well-designed system the line driver will be the limiting factor.

To quantify this, consider the simplified circuit of figure 40a, where the magnetisation inductance L_M has been neglected. This allows to consolidate the (undesired) coupling impedances on both sides of the transformer in a lumped impedance on the primary side.

8.1 Theory — Class B amplifiers

The power dissipated inside the amplifier in figure 40a can be calculated easily if we limit ourselves to a sinusoidal signal and resistive loading. To simplify the calculations a symmetric supply is assumed, but the results are equally valid for an AC-coupled single-supply amplifier.

The amplifier load Z_L can be represented by a single impedance[59]

$$Z_L = (Z_M + Z_{CP})\, N^2 \tag{2}$$

as shown in (figure 40b). The power drawn from the supply is found from $P_S(t) = {}^1\!/_2 V_S I_S(t)$. Averaged over a signal period, the supply current equals $\overline{I_S(t)} = \frac{1}{\pi}\int_0^\pi I_S(t)\mathrm{d}t = I_L\frac{2\sqrt{2}}{\pi}$. The power balance is then easily found from

$$\begin{aligned} P_L &= V_L \cdot I_L \\ P_S &= \frac{V_S}{2} I_L \frac{2\sqrt{2}}{\pi} + P_Q \\ P_A &= P_S - P_L \end{aligned}$$

Knowing that $V_L = Z_L I_L$ the dissipation inside the amplifier can then be written as

$$\begin{aligned} P_A &= P_Q + V_S\frac{\sqrt{2}}{\pi}\frac{V_L}{Z_L} - \frac{V_L^2}{Z_L} \qquad (3) \\ &= P_Q + V_S\frac{\sqrt{2}}{\pi} I_L - I_L^2 Z_L \end{aligned}$$

For an ideal amplifier $P_Q = 0$ and $0 \le \sqrt{2}V_L \le {}^1\!/_2 V_S - V_O$.

For a fixed load the driver dissipation is non-monotonic and reaches a maximum when the amplitude is $V_L = \frac{1}{\sqrt{2}\pi} V_S$ (figure 41). Thus, maximal efficiency is reached at maximal signal amplitude. The theoretical efficiency at this point (with $V_O = 0$ and $P_Q = 0$) is equal to $\frac{\pi}{4}$ or about 78.5%.

8.2 Safe operating area

The temperature of the silicon junction inside the line driver T_J must not be allowed to exceed the maximum

[59] Refer to table 4 on page 39, «Nomenclature», for a list of symbols.

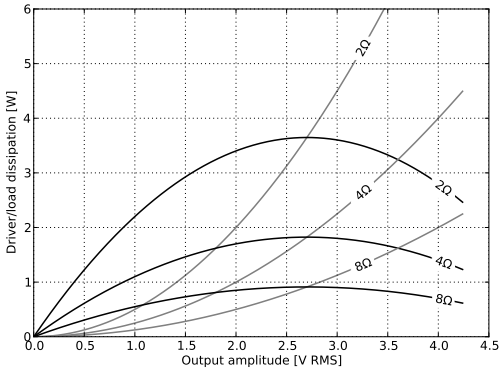

Figure 41: Ideal class B amplifier with ±6 V power supply and a sinusoidal signal: dissipation in the load (grey) and the amplifier for three loads as a function of signal amplitude.

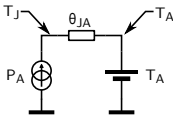

Figure 42: Thermal model used for steady state analysis.

specified in the line driver datasheet. Otherwise, a substantial reduction in life time will result.

T_J can be calculated from the heat generated inside the amplifier (P_A), the ambient temperature and a thermal model. With the simplest possible thermal model (shown in figure 42) consisting of only the junction-to-ambient thermal resistance θ_{JA}[60]

$$T_J = T_A + P_A \theta_{JA} \tag{4}$$

In practice, θ_{JA} is usually an input for the design flow, not a design requirement[61]. The designer specifies the maximal allowable T_J, the achieved θ_{JA} and the maximal ambient temperature for which correct operation must be guaranteed. With these and eq. 4, the maximal power dissipation can be derived.

[60][Sto06] discusses how to interpret thermal resistances.

[61]The designer of the printed circuit board (PCB) will strive for a minimal thermal resistance, given the constraints on the PCB. While adding a heatsink or using a four-layer PCB instead of two layers can substantially decrease θ_{JA}, but such changes are usually not possible given the additional cost.

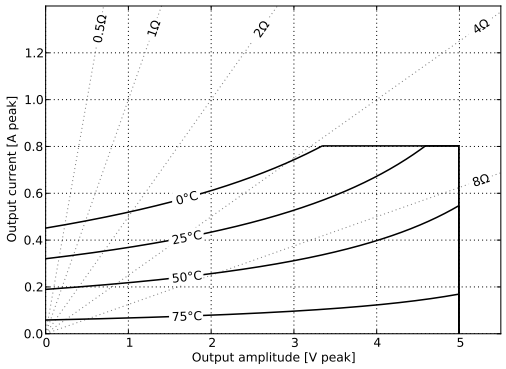

Figure 43: Safe operating area (SOA) of the output amplifier for various ambient temperatures assuming $\theta_{JA} = 50$ K/W (bottom left corner is safe).

The operating area of a class-B amplifier is not representable in a single graph as it depends on load, supply voltage and cooling.

Equation 3 allows to calculate the derived $P_{A,\max}$ back to V_L and I_L values that are do not result in excessive heating.

A typical safe operating area (SOA) is shown in figure 43. Any operating point inside is acceptable. The area is bounded by three limits. As discussed, thermal aspects limit the maximal power dissipation. Beyond these limits the thermal protection of the line driver will trigger. Secondly, the maximal current (800 mA RMS) is seen in the flat top line. The maximal output voltage (limited by V_{OH} and V_{OL}) results in the straight line on the right[62].

Obviously the exact limits will depends on the line driver. Figure 43 shows the SOA for a typical NCS5651-based design with a typical $\theta_{JA} = 50$ K/W.

Although voltage-versus-current is the normal representation of safe operating area, the line driver can only control one of these variables: voltage and current are linked through the load impedance. Figure 44 displays exactly the same information as figure 43 but might be easier to work with. Constant current values are now represented as canted lines.

So far we have only considered the V_L and Z_L, i.e. the voltage and load as seen by the amplifier. In prac-

[62]The maximal current, V_{OL} and V_{OH} are largely temperature-indepedent; therefore in figure 43 T_A influences only the maximal power dissipation inside the line driver.

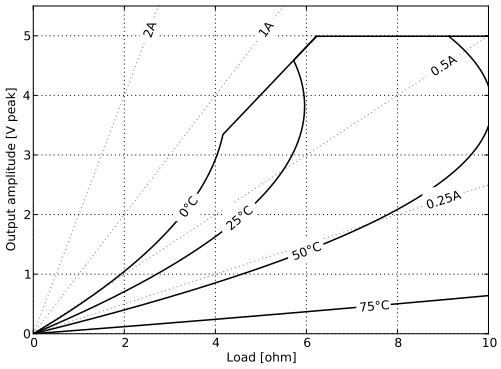

Figure 44: Alternative representation of the same SOA shown in figure 43 (bottom right corner is safe).

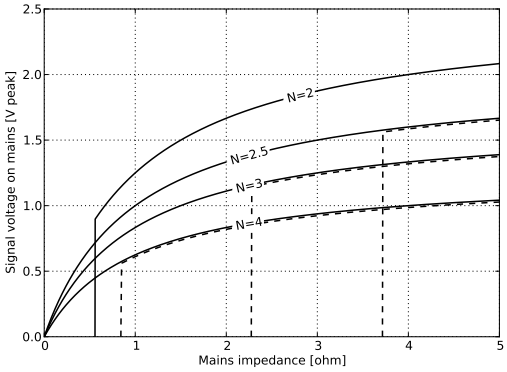

Figure 45: Voltage V_M on the mains as a function of mains impedance for four ratios N and two temperatures: 25°C (solid lines) and 75°C (dashed lines).

tice, these are different from V_M and Z_M due to the transformer. Obviously, a higher transformer ratio N results in a lower mains amplitude:

$$V_M = \frac{Z_M}{Z_M + Z_{\text{CP}}} \frac{1}{N} V_L$$

At the same time, a high N is desirable to drive low mains impedances. From equations 2 and 3

$$\begin{aligned} \min\{Z_M\} &= \frac{1}{N^2} \min\{Z_L\} - Z_{\text{CP}} \\ &= \frac{1}{N^2} \underbrace{\frac{\max\{P_A\} - P_Q}{V_S \frac{\sqrt{2}}{\pi} - V_L}}_{\max I_L} \frac{1}{V_L} - Z_{\text{CP}} \end{aligned}$$

Adding the current limit I_{lim} this becomes

$$\begin{aligned} \min\{Z_M\} &= \frac{1}{N^2} \max\left[\frac{\max\{P_A\} - P_Q}{V_S \frac{\sqrt{2}}{\pi} - V_L}, I_{\text{lim}} \right] \frac{1}{V_L} \\ &\quad - Z_{\text{CP}} \end{aligned}$$

Thus, lower transformer ratios are will perform better for higher mains impedances and vice versa. Figure 45 shows an example for $Z_{\text{CP}} = 1\,\Omega$.

8.3 Design considerations

Figure 3 shows that V_L is preferably chosen as high as possible to maximise the efficiency, i.e. $V_L = {}^1\!/_2 V_S - V_O$.

The coupling impedance Z_{CP} is usually designed to be as low as possible, and is therefore an input to the design.

The maximal ambient temperature T_A is usually a design requirement. Through equation 4, this will fix $\max\{P_A\}$[63].

As a result, the designer can only choose the transformer ratio N to optimise the safe operating area. By defining the minimal Z_M for which correct operation must be guaranteed, N can be derived. Of course, the difficult part is selecting $\min\{Z_M\}$; the scientific literature may help, or measurement campaigns may be required[64].

9 Power supplies

9.1 Typical consumption

The modem itself has two 3.3 V supply pins, analogue and digital[65]. Typically they are connected to the same rail, but individual decoupling capacitors are required. The exact current drawn will depend on the

[63] Note that the maximal allowable T_A will also depend on the enclosure design.

[64] The mains impedance is quite complex, varying with place and time. For more information, refer to e.g. [CK08].

[65] The modem also has a 1.8 V rail which is fed from an internal linear regular. The user must only provide a 1 µF capacitor on the board to ensure stability of the regulator. No other connection should be mad to this pin.

firmware; 35 mA is a typical value. Any current delivered to external loads (for instance, LEDs) must be added to this.

The line driver, whether standalone (NCS5651) or embedded in the NCN49599, also has a 3.3 V pin to interface with the modem. It should be powered from the same source as the digital pin of the modem. The power dissipated in the IC itself is typically negligible, but as with the modem LEDs &c. must be taken into account.

Of course, substantial currents are drawn from the main supply of the line driver (typically 12 V). The exact value will depend on the mains impedance, the coupling circuit — especially the transformer ratio — and the signal amplitude. A peak current of 800 mA is a typical value.

9.2 Modem requirements

The 3.3 V rail is used to power the analogue circuitry: filtering in the modem, analogue-to-digital conversion and usually external filtering. Noise and interference on this rail can interfere with the demodulation. Because the received PLC signal is often exceedingly weak, it is easily swamped by coupled supply noise.

The 3.3 V regulator must therefore be designed carefully. With care a switching regulator can be used, but it is advisable to select a part with a switching frequency substantially above the highest carrier frequency — 500–1000 kHz is ideal. Most modern switching regulators, whether AC/DC or DC/DC, reduce the switching frequency under low load to improve efficiency (through cycle skipping or frequency foldback). The actual switching frequency in receive mode must be verified in the final design as it will usually be dependent on the precise load. In transmit mode, the noise on the supply rail is less critical.

The layout of the regulator circuit is critical to keep noise under control. The regulator datasheet usually provides detailed guidelines.

9.3 Topology

9.4 Input filter

In modern electronics, AC/DC supplies are virtually always realised as switch-mode power supplies (SMPS) due to their efficiency, size, weight and cost..

In PLC systems their application demands careful attention.

9.5 Output filter

10 Modem control and interface

Typically, an external microcontroller interfaces with the modem over the universal asynchronous receiver and transmitter (UART), described in [ON 14a, pp. 24–25]. The baud rate of the UART is set by the pins BR0 and BR1. Changing the baud rate during operation is rarely required; connecting these pins to ground or supply is recommended to simply the layout. The baud rate is usually set as high as possible (115200 baud) but if the UART signal have to pass opto-couplers or other slow components a lower baud rate might be appropiate.

The firmware of the modem must be loaded following every modem reset.

This may done over the UART (described in [ON 14a, "Boot loader"]. Alternatively the modem can retrieve the firmware autonomously from an external flash memory attached to the SPI port. The first option is usually less expensive but requires about 20 KiB programmable memory in the microcontroller.

It is strongly recommended to connect unused input pins to ground. As with any CMOS logic, the voltage on a unconnected input pin can float; the input impedance of the NMOS-PMOS transistor pair used in a CMOS input cell is extremely high. If the voltage drifts away from ground or the supply, both NMOS and PMOS transistors will conduct somewhat, leading to increased power consumption. Unexpected behaviour can result.

In addition, floating input pins increase the sensitivity to electromagnetic interference (EMI).

The input pins of the JTAG port (TID, TCK, MTS, TRSTb) are not used in an application design; connect them to ground to ground or 3.3 V. The output pin TDO must be left unconnected.

If the GPIOs IO4–IO10 and the SPI input pin SDI are not used, these must be tied too. If future use as an input or output is possible, a pull-up or pull-down resistor may be preferable to a direct connection to a rail. In the latter case, care must be taken to ensure the pins are always configured as inputs.

Connecting not-connected (NC) pins to ground is also strongly recommended.

11 Galvanic isolation

Substantial savings can be realised by switching to a non-isolated design. The transformer can now be removed, but two significant disadvantages arise.

The transformer not only isolates, but also increases the impedance seen by the power amplifier. A transformer-less design will show good performance on a highly impedant mains; at lower impedances (common in offices and dwellings) the maximal transmission amplitude will be lower. In turn this results in a lower communication range.

Additionally, a non-isolated modem printed circuit board is no longer touch-safe. If this is a requirement, the user microcontroller section and the modem section can be isolated at the cost of three isolators (for reset, TxD and RxD) and an additional power supply.

Whether the optimal design is isolated or non-isolated strongly depends on the application requirements. Your ON Semiconductor representative can help you and provide more information.

12 Filters

The filters on the evaluation daughterboard are designed for a particular carrier frequency set (refer to section 19 for the responses). This section describes how to redesign the filter for a specific application.

12.1 Receive filter

A well-designed receive filter is critical to improve the reception performance of the modem.

The exact requirements of the receive filter depend on the application.

The optimisation goals relate to the suppression of out-of-band noise and the carrier frequencies (gain mismatch and settling time). To optimise, we can tweak the gain, the filter order and the pole-zero placement. We will discuss these in this order.

Out-of-band noise is easy — ideally it would be suppressed completely. To understand this consider figure 46. The modem digital mixer hardware is capable of rejecting out-of-band interference very well. However, the ADC still needs to handle the entire input signal, including interference. To improve the dynamic range an amplifier with auto-gain control (AGC) precedes the DAC. High interference will trigger a low AGC gain. In turn, this will degrade

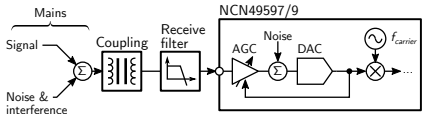

Figure 46: Influence of external interference on the demodulation.

the input-refered quantisation noise and therefore the signal-to-noise ratio (SNR) of the digitised signal. As a result the robustness of the demodulator will improve.

Suppressing out-of-band interference *before* the signal enters the modem allows a higher gain in the AGC block without overloading the ADC.

In addition the ideal filter has a low carrier gain mismatch, i.e. keeps the gain difference for all carrier frequencies below a tolerance.

The final optimisation goal is the settling time at the carrier frequencies. Settling time will result in inter-symbol interference. It this becomes too high it will degrade the demodulation, leading to bit errors.

The choice of the carrier signal gain depends on the environment. If a high signal attenuation is expected a high gain is useful. However, clipping caused by interfering signals must be avoided, giving an upper limit for the gain.

When in doubt, choose a lower gain. A reasonable starting value is -6 dB.

Most commonly a band-pass filter will be chosen. The width of the pass band depends on the carrier frequencies.

Once the filter order and the tolerable gain mismatch and group delay are chosen the exact filter cut-off frequencies can be derived. Lower acceptable gain mismatch usually implies a wider filter. A Butterworth filter is recommended as it provided a good balance between dispersion and quick falloff.

To achieve a low gain mismatch with sharp filters, the cut-off frequency of the low-pass and highpass filters are about the same as the lowest resp. highest expected carrier frequency.

ON Semiconductor modems provide a low-noise operational amplifier for filtering the receive signal.[66] This opamp can implement two of the poles of the designed filter. The filter response can be configured with external components.[67] If the design includes

[66] Refer to [ON 14a, p. 24] or [ON 14b] for details.

[67] To calculate the values of the passive components a fil-

more poles, an external opamp or passive filtering circuit must be added.

Once the filter has been designed its sensitivity to component tolerances must be evaluated. If the sensitivity it too high, components with a tighter tolerance must be selected or the filter must be redesigned.

12.2 Transmit filter

The transmit filter is located between the DAC inside the modem and the power amplifier. The purpose of the transmission filter is to filter the harmonics and spurious tones inevitably generated by the modem transmission DAC to standard-complaint levels. This immediately implies the applicable standard will drive the design of the transmission filter.

Two widely used standards are CENELEC EN 50065-7 ([EN01a]) and the FCC regulations, but many others exist.

To increase suppression of the carrier harmonics, the cut-off frequency of the transmit filter is usually close to the highest carrier frequency. As a result, the filter gain drop will already be pronounced at the carrier frequencies. In addition, both carrier frequencies will see a different gain. Thus the amplitude for the mark and space bits will be slightly different. As long as this gain mismatch remains below a few dB it has no ill effect. The mismatch can be reduced, but only at the cost of lower suppression of harmonics.

13 Cost reduction

The evaluation kits are designed for flexibility and good performance within a reasonable cost. A lower cost version is possible, but the designer should carefully weigh the trade-offs involved.

- A significant reduction in component cost can be achieved by replacing L1[68] by a 1 Ω resistor. However, this reduces the available output voltage if the mains has a low impedance at the transmit frequencies. L1, D4, D5 and C17 can be removed altogether when supplying the power amplifier from 6 V or less. Obviously, the output swing and communication range is curtailed in such a design.
- If CENELEC compliance is not targeted, the transmit filter may be reduced to a third order filter; even a first-order filter is possible if the resulting severe distortion can be accepted.
- The transmit filter on J daughterboards has been tuned for multichannel PL110. If single-channel is sufficient, the filter order may be reduced to three without exceeding CENELEC distortion limits.
- When operating in environments with low interference, the lowpass section in the receive path may be removed. We strongly advise extensive testing to ascertain communication quality if this option is chosen.
- As explained in section 2.4, R36 is required for KNX PL110 compliance. If this is not a requirement, the resistor may be shorted out; IO3 must then be removed from the net Tx_Eb.
- The filter at the input of the SMPS can be reduced to a second order filter, instead of the current fourth order filter. However, conducted emission will be increased and a careful evaluation is required. Additionally, it might be possible to eliminate L12 if the SMPS ripple and noise is small enough.
- The ON-PL110 embedded software does not need zero-crossing detection. The zero-crossing circuit is included on the motherboard solely to support other software solutions.

14 Multi-phase designs

The evaluation kit is a single-phase design. However, for repeaters and concentrator multi-phase designs, especially three-phase designs, are usually required. These designs are fairly similar to the single-phase case but a few aspects must be considered.

It is important to remember that concentrators and to a lesser extent repeaters play vital roles in the PLC network. If these central node perform badly, the entire network will under-perform. Thus, excessive cost-cutting in the multi-phase designs is counterproductive.

ter design program such as Texas Instruments FilterPro or Microchip FilterLab may be employed. Both programs can be obtained free of charge.

[68] In this section, component designators refer to the components in the daughterboard schematic.

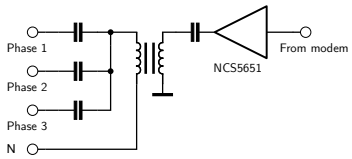

Figure 47: Driving three phases with a single amplifier.

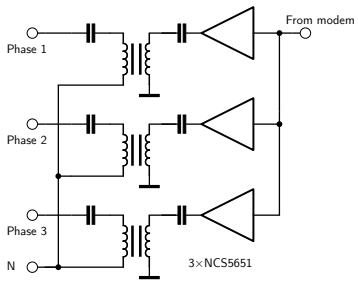

Figure 48: Driving three phases with dedicated amplifiers.

The power supply can be derived from one phase, as in the evaluation kit. This is the most simple and economical solution. As an alternative, the power supply can be derived from all phases. This would allow the modem to continue to transmit and receive even if the mains voltage on some phases disappears. Whether this circumstance is realistic depends on the end application.

14.1 Transmission path

It is possible to drive all phases with a single NCS5651 opamp (figure 47). However, this opamp will generally see a very low impedance, as it will have to drive multiple phases in parallel. In addition, concentrators and repeaters commonly see a very low impedance when multiple users are connected. Therefore, the recommended solution is to use an NCS5651 opamp per phase (figure 48).

This topology also improves the tolerance to fault conditions. For instance, if the impedance on one phase is so low that the line driver overheats and thermal shutdown is triggered, the drivers on other phases will continue to operate normally; communication with nodes connected to these phases will still be possible.

Monitoring all line drivers for over-current protection and thermal shutdown is recommended.

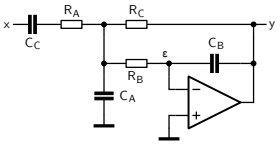

Figure 49: First stage of the recommended active receive filter.

It is recommended to use an NCN49597 modem with multiple NCS5651 amplifiers. Using of an NCN49599 in order to save an external NCS5651 is possible, but not recommended. The NCN49599 is primarily designed for minimal cost, thus it has no TSDflag pin. This precludes the application microcontroller from reacting to thermal shutdown.

The transmit filter can be adopted from single-phase designs without changes.

14.2 Reception path

Regardless of the coupling topology chosen, it is strongly recommended that the receive path is sufficient filtered to suppress interference.

Reception on multiple phases simultaneously is inherent in the shared-transformer coupling shown in figure 47. The single-phase receive circuit of the evaluation boards can be adopted without changes. The same applies to non-isolated coupling circuits.

When separate coupling transformers are used, the receive circuit must be modified.

The design may be started from the recommended receive filter, a two-stage active filter.

The first stage of this filter uses a low-pass multiple-feedback (MFB) topology with AC-coupling (figure 49).

The topology can be modified for multiple-phase designs by adding input branches (figure 50).

Note the values of the passive components must be adjusted. Denoting the values of the single-phase design as R_A, R_B &c. and those of the three-phase as $\widehat{R}_A$, $\widehat{R}_B$ &c., a reasonable starting point is $\widehat{R}_A = R_A$ to keep the input impedance identical. Choosing $\widehat{C}_C = C_C$ keeps the high-pass cutoff frequency stable[69].

Unfortunately the optimal gain is not obvious. To retain the same signal level, $\widehat{R}_C$ must equal R_C, but this will increase the noise and interference level as

[69] Note the design procedure of the single-phase filter assumes C_C has no influence on the low-pass filter poles. The assumption holds well enough for reasonable design inputs.

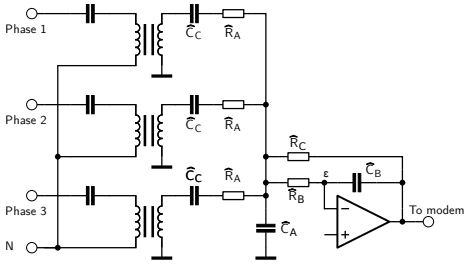

Figure 50: Receiving on three phases simultaneously with separate coupling transformers.

three phases are now contributing. The increase in noise floor depends on the noise correlation between the three phases.

In one extreme, all noise is perfectly correlated. In this case, the noise *amplitude* will be three times higher (10 dB). Under these conditions a lower gain will be appropriate to avoid swamping the receiver. If all noise is uncorrelated, the noise *energy* will be three times higher (5 dB).

Field measurements in the expected environment can help determine the typical noise.

As a starting point, a gain of half the gain of the single-phase design is a reasonable compromise ($\widehat{R}_C = {}^1\!/_2 R_C$).

With R_A and R_C chosen, no degree of freedom remain.

To retain the low-pass pole placement, the gain from the output of the amplifier (denoted y) to the net ϵ must remain the same. For the single-phase topology this gain at DC is given by

$$H_{y\to\epsilon}(s=0) = \frac{R_B}{R_B + R_A||R_C}$$

With three input branches R_A must be replaced by $\widehat{R}_A/3$. From the requirement $\widehat{H}_{y\to\epsilon}(s=0) = H_{y\to\epsilon}(s=0)$ the value of $\widehat{R}_B$ follows:

$$\widehat{R}_B = \frac{\widehat{Z}}{1 - \frac{\widehat{Z}}{\widehat{R}_A/3}} \text{ with } \widehat{Z} = \frac{\widehat{H}_{y\to\epsilon}(s=0)\widehat{R}_C}{1 - \widehat{H}_{y\to\epsilon}(s=0)}$$

Following immediately from $\widehat{R}_B$, $\widehat{C}_B = {}^{C_B R_B}\!/_{\widehat{R}_B}$.

The time-constant of the impedance level at ϵ must also remain the same when changing to a three-phase topology:

$$\tau_\epsilon = C_A\,(R_A||R_B||R_C)$$

Thus,

$$\widehat{C}_A = C_A \frac{R_A||R_B||R_C}{\widehat{R}_A/3||\widehat{R}_B||\widehat{R}_C}$$

For more information and reference schematics for a three-phase design, contact your sales representative.

15 Layout

15.1 Electrical aspects

The layout of PLC modem boards is critical because the received signal levels are very low. At the same time the switching power supply, an important sources of interference and noise, is located close to the sensitive analogue receiver. It is recommended to maintain an uninterrupted plane for modem ground as much as possible. As an example, the ground planes of a motherboard and an NCN49597 daughterboard are shown in figures 59 and 60.

Specific information on the layout of the protection circuit appears in section 7.4.

15.2 Thermal aspects

Thermal aspects must also be considered during the layout. General recommendation on thermal PCB design with QFN packages may be found in [Amk08, ON 02]; application note [Sto14] provides specific information for the NCS5651.

The following paragraphs give some additional hints.

For illustration, a cut-through of a two-layer board with a QFN package line driver is shown in figure 51. Copper is shown dark grey; the solder masks are shown hatched; the solder is shown as light grey.

To keep the silicon temperature of the NCS5651 or NCN49599 line driver acceptable during operation, cooling must be provided. To this end, the QFN package is provided with a central *exposed pad* on the bottom side (figure 51, 1). In a typical design, this exposed pad carries the bulk of the heat flow from the silicon die to the ambient. As a result, the design of the layout around the exposed pad is critical for the thermal performance of the system.

The exposed pad should be soldered to a single solid copper pad on top copper, with thermal vias connecting it to the bottom copper[70].

[70] Only two vias are shown in figures 51–56 (annotated as 4), but a more substantial number is usually required.

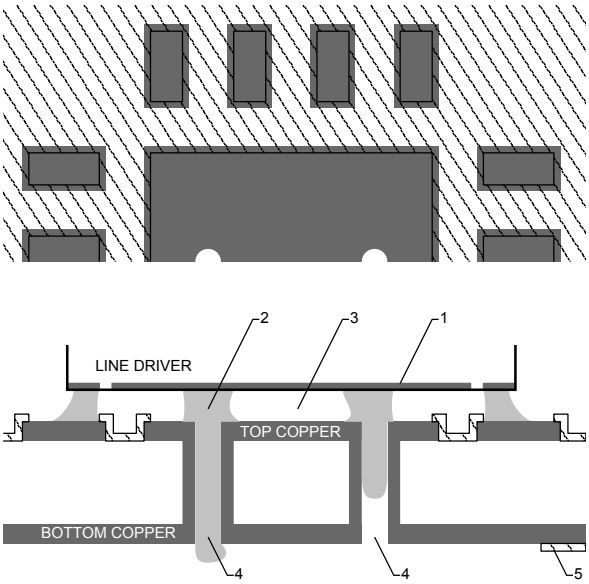

Figure 51: Top view and cut-through of a line driver (NCN49599 or NCS5651) soldered on a two-layer board: no solder mask on the thermal vias. Numbers are referred to in the text.

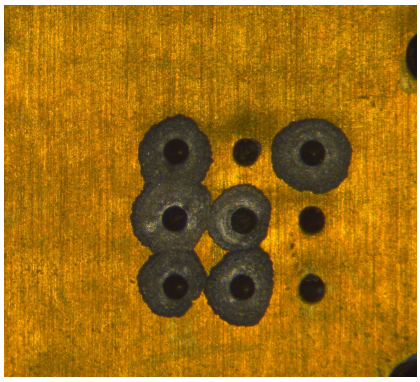

Figure 52: Result of solder paste wicking through thermal vias during reflow. Photograph of board 1, bottom side.

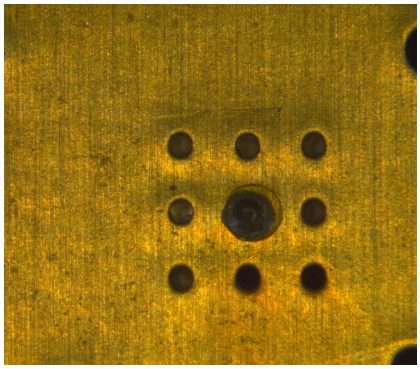

Figure 53: Result of solder paste wicking through thermal vias during reflow. Photograph of board 2, bottom side.

However, during reflow the solder paste deposited on the central pad wicks can wick into the thermal vias (figure 51, 2), pulled down by gravity. Often the solder flows out to the bottom side, as is evident in figures 52–53 showing two board from the same batch[71].

Wicking causes voids (3) below the exposed pad, dramatically increasing the thermal resistance.

To reduce wicking the lowest possible via diameter should always be used. Usually this is not sufficient. Filled vias can completely stop wicking. In addition, if copper-filled vias are used, the heat conduction is much improved. However, filled vias are expensive.

An alternative is to *tent* the thermal vias[72], i.e., cover them with solder mask (6). By the vias on the top side (or both top and bottom side), solder wicking is much reduced (figures 54 and 55). Tenting the vias on the bottom side only should be avoided, as the solder paste outgassing during reflow can cause voids (figure 56, refer to [Sto13] for details).

Adding more vias improves the thermal resistance from top copper to bottom copper, but it also increases the solder mask covering. Providing 7–12 vias is recommended for the NCS5651; for the NCN49599, 10–16 vias. Via plating should be as thick as possible.

For the NCN49599 the copper area for the exposed pad is quite large. To force even solder paste distribution on the pad it is recommended to provide thin lines in the stencil below the QFN, forming a 2×2 or 3×3 pattern (refer to [ON 02, figure 5] for details).

[71] Comparing both photographs, it is clear that even within a single batch the amount of wicking is unpredicable. As a result, increasing the amount of solder paste deposited in the hope of compensating the lost solder rather tricky, especially for smaller production runs.

[72] Obviously the size of the solder mask cover should be as small as possible to avoid degrading the thermal resistance between the exposed pad and the top copper. Usually the minimal overlap between copper and solder mask close to a hole is defined in the design rules of the printed circuit board manufacturer. This design rule stems from the inevitable misalignment between the solder mask and the copper. The latter misalignment is clearly seen in figure 55.

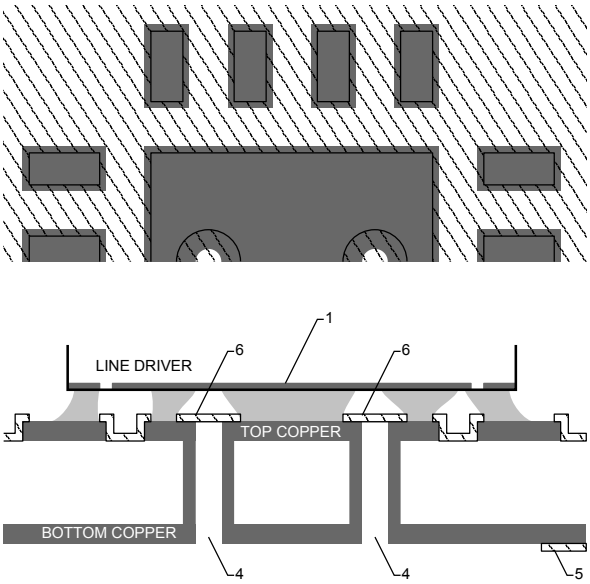

Figure 54: Top view and cut-through of a soldered line driver: solder mask tenting on top of the thermal vias.

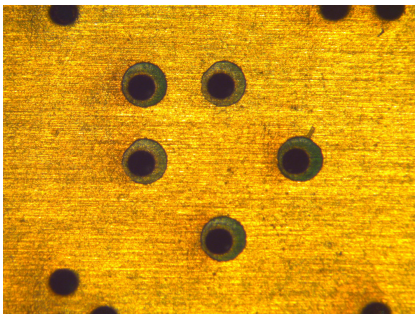

Figure 55: Photograph (bottom side) after reflow of a board with the thermal vias tented on both sides.

This stencil pattern also improves venting of the gases produced during reflow ([Sto13, p. 4]).

Improving the heat transfer from the thermal vias to the ambient must also be considered.

If possible,[73] select a thin printed circuit board and thick copper.

If two-sided assembly is used, added solder paste on the bottom side (figure 56, 7) improves the heat spreading considerably.

[73] Usually the designer is constrained by other consideration; especially for QFNs fine copper tracks are required and therefore thin copper foil.

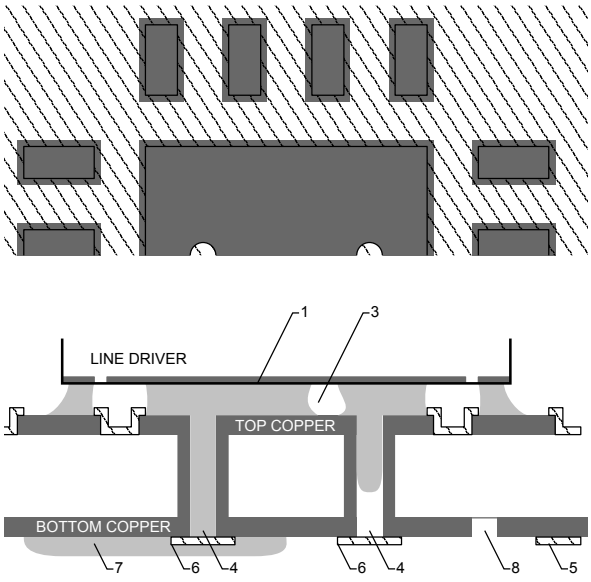

Figure 56: Top view and cut-through of a soldered line driver: solder mask tenting below the thermal vias.

Remember that the solder mask adds a small but notable thermal resistance. If gold-plating is used, keep the bottom solder mask as open as possible (5) to improve heat transfer from the copper to the ambient air. If you don't use gold-plating, keeping the soldermask is advisable to avoid oxidation and a increasing thermal resistance over time.

Keep the all-important bottom copper plane as intact as possible, especially close by the thermal vias. Avoid breaks (figure 56, 8) as much as possible. When tracks must be routed in the ground plane, they should be kept as short as possible. Multiple short bottom are preferred to a single long one (figure 57c–d).

To allow the heat to spread as quickly as possible from the thermal vias, tracks in the bottom copper layer should be aligned to the heat flow as much as possible, i.e. pointing towards the thermal vias (figure 57a–b).

The optimal design is dependent on the printed circuit board technology. Therefore, the designed should ask the board manufacturer for design guidance. In practice, a trial run is usually required to verify correct soldering. Röntgen photography is very helpful as it can show the exact solder density and the voids below the exposed pad.

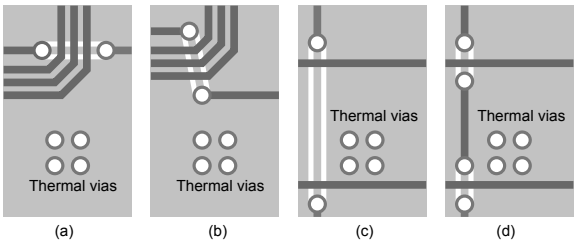

Figure 57: Keeping the ground plane intact: tracks in bottom copper aligned to the heat flow (b) are preferred to tracks an orthogonal orientation (a); multiple short tracks (d) are preferred to a single long track (c).

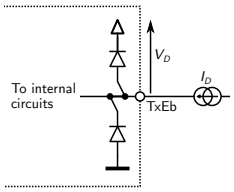

Figure 58: Using an ESD protection diode to estimate the die temperature.

On a four-layer board it is recommended to use the inner layers are used in addition to, not as a replacement for, the bottom copper.

Measuring the thermal resistance

Predicting and controlling the cooling performance of a line driver designis difficult. Thus, measurements are useful both during development and production. By measuring the temperature of the line driver die T_J, the ambient temperature T_A and the power dissipated in the line driver P_A, the thermal resistance can be easily calculated:

$$R_{JA} = \frac{T_J - T_A}{P_A}$$

T_A and P_A are easily measured, in contract to the die temperature.

A good estimate for T_J can be derived by measuring the temperature of an on-die ESD protection diode. Pins of A a typical protection circuit is shown in figure 58 (left).

By sourcing (as in figure 58) or sinking (not shown) a small current on a pin, one of the protection diodes can be made to conduct. The temperature of the diode follows from the Shockley diode equation

$$I_D = I_{D,S}\left(e^{V_D/(nV_T)} - 1\right)$$

and substituting $V_T = {}^{kT}/_{q}$

$$T = \frac{qV_D}{\ln\left(\frac{I_D}{I_{D,S}} + 1\right) nk}$$

The reverse bias saturation current I_S can be estimated from a measurement with the device switched off. This ensures that the die temperature closely approximates room temperature. In practice, calibrating for multiple temperature is a good idea[74].

The current source can be conveniently constructed from a voltage source (~25 V) and a resistor[75].

V_D and I_D can then be measured with a bench multimeter. Care must be taken to minimise cabling capacitance and resistance [ON 11].

16 Troubleshooting

We describe some common problems, together with suggestions to debug the problem. Contact your sales representative or field application engineer for further help.

Also refer to section 4

16.1 The modem is not booting (firmware upload over the serial interface)

Problem: you try to upload a binary image over the serial interface (typically from a microcontroller) and the modem does not boot correctly.

To debug, verify that the baudrate is correctly set (BR0 and BR1 pins set according to [ON 14a, table 10]). Refering to [ON 14a, figure 3] for the correct boot sequence, observe the signal on the pins IO2, RESB, TxD and RxD with an oscilloscope. Ensure

[74] More advanced measurement techniques exist that eliminate any calibration step, and minimise the parasitic effect of cabling. For more information, refer to [ON 11].

ON Semiconductor offers a range of integrated circuits for diode temperature measurements based on these techniques.

[75] Of course, this will not result in a truly constant current variations in V_D will influence I_D. However, for all practical purposes the accuracy is satisfactory if the source voltage is substantially bigger than V_D.

SEN and IO2 are low before RESB is released (i.e. rises) and remains so for the entire download.

The modem should respond with an STX byte (02_H) on the TxD line.

- If the STX byte is not transmitted, ensure the modem is clocked at the correct frequency. The easiest way to do so is to allow the internal code to start (IO2 high, SEN low, RESB pulled low and then released). Once started, the boot loader indicates it is running by inverting IO0 every second (i.e. the period should about 2 s). Alternatively, observe the XOUT pin with a low-capacitance probe. A small 48 MHz sinusoidal signal is expected.
  - If the modem clock is not correct, verify the crystal selection guidelines specified in [ON 14a, section 2.2] are followed. Check the EXT_CLK_E pin is pulled low to enable the internal crystal oscillator.
  - If the modem clock is correct, ____.
- If the STX byte is transmitted, use it to measure the baud rate. Then, verify that the download starts correctly as illustrated in [ON 14a, figure 3] by observing RxD. Of course, the same baud rate should be used.
  - If the modem does not respond at the end of the download sequence, ensure the correct length was transmitted.
  - If the modem responds with a 15_H byte (NAK), check the timing constraints noted in [ON 14a, figure 3] and [ON 14a, section 5].
    Also ensure an image with a correct checksum appended was transmitted. In firmware images prepared by ON Semiconductor, «CRC» appears in the filename if a checksum was appended.

  If the modem responds with an ACK byte, the firmware was loaded correctly. If you do not observe the expected behaviour, again look at IO0. Firmware prepared by ON Semiconductor will indicate its status on this pin. Both the IEC and the ON-PL110 firmware variants will — if unconfigured — emit a 1 Hz signal (0.5 s low, 0.5 s high).

16.2 The output amplitude is off

If the measured amplitude of the signal is not matching the expected value, these hints might be useful.

- Ensure the measurement is done correctly. A common error is to set the sample rate too low. Refer to section 4.1.6 on page 16 (esp. figure 22) for more information.
- Ensure both carrier frequencies are correct.
- Measure the amplitude for *both* carrier frequencies. Due to filtering, they will be slightly different. The difference can give a clue to the source of the problem.
- Measure the amplidues at every amplifier output through the signal chain. If a recommended active filter is used, measure at the output of the modem, the output of the external opamp (not applicable for all filter variants), and the outputs of both the «A» and the «B» amplifier of the NCS5651 line driver.[76]
- Ensure the amplitude at the secondary side of the transformer[77] is about equal to the amplitude at the output of the line driver (i.e. before and after the low-voltage decoupling capacitor).
- Calculate the expected gain for each carrier frequency for each filtering stage. Verify this against the measured amplitudes.

17 Standard-compliance

A wide range of standards is applicable to power line communication systems. Which standards are required depends on the application. For instance, ON-PL110-based application will disregard IEC 61334. Commonly encountered standards are

IEC 50065 describes the hardware requirements of power line communication equipment with carrier frequencies up to 148.5 kHz. Originally a European standard, it has been widely adopted, though sometimes with modifications.
The standard is focused on ensuring coexistence amongst PLC networks.

[76] On the standard evaluation kit daugherboards, the relevant test points are `MTx`, `Tx`, `OutA` and `PL`, respectively.

[77] With a standard evaluation, it is most convenient to measure this on the motherboard.

IEC 50065-1 prescribes the maximal (unintended) conducted emissions.

IEC 50065-7 prescribes the minimal impedance at the mains connector. A sufficiently high impedance avoids disturbing other networks.

IEC 60950-1 is the applicable safety standard for most application.

IEC 61000-4 prescribes the conducted electromagnetic interference a design must tolerate.

IEC 61010-031 strives for safe measurement proceduces. Although not applicable to the final installation, it must be considered to ensure safety during development, field tests and installation.

IEC 61334 is a protocol commonly used for metering. Often the protocol is conveniently if lazily called «IEC», a convention followed in the present document.

IEC 61334-5-1 describes the lower layers.

IEC 61334-4-32 describes the link layer control (LLC) part of the IEC 61334 standard.

Linky is a protocol building on IEC 61334, adding some features. It was designed by the French distribution grid operator ERDF for the French smart metering rollout.

Symbol	Unit	Description
V_L	V RMS	Voltage at the line driver output
I_L	A RMS	Current out of the line driver output
Z_L	Ω	Load impedance at the line driver output
P_L	W	Power dissipated in the load
V_S	V DC	Line driver supply voltage
I_S	A peak	Line driver supply current
P_S	W	Line driver consumed power (equal to $P_L + P_A$)
P_Q	W	Line driver quiescent dissipation
P_A	W	Total line driver dissipation
V_M	V RMS	PLC signal voltage on mains
Z_M	Ω	Mains impedance at the PLC carrier frequency
V_O	V	Saturation voltage to rail[79]
T_J	K	Junction temperature
T_C	K	Case temperature top
T_P	K	Case temperature bottom
T_{Cu}	K	
T_A	K	Ambient temperature
θ_{JA}	K/W	Thermal resistance between junction and ambient
!		The notation of other thermal resistances follow this convention.

Table 4: Nomenclature

Part III
Appendix

18 Nomenclature

A list of symbols used appears in table 4[78].

19 Evaluation kit design

The copper design of the revision 6 motherboard and the revision 3 NCN49597 daughterboard are shown in figures 59 and 60[80]. The bottom copper is shown mirrored, i.e. as if seen through the board.

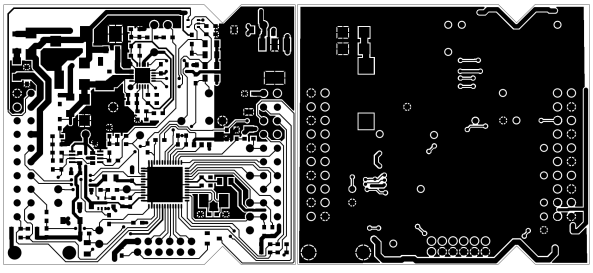

Figure 59: Layout of the rev. 3 NCN49597 daughterboard: top copper (left) and bottom copper.

[78] [Sto06] discusses how to interpret thermal resistances.

[79] The voltage drop to the rail is a function on the current. However, in this document we neglect this dependency and assume V_O to be constant to simplify calculations.

[80] Figures 59 and 60 are not to scale.

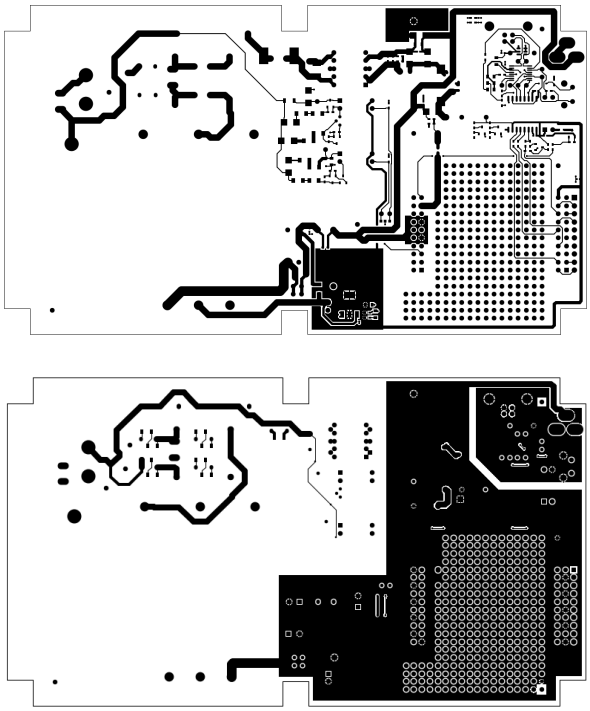

Figure 60: Layout of the rev. 6 motherboard: top copper (top) and bottom copper.

References

[Amk08] Amkor Technology. Application notes for surface mount assembly of Amkor's MicroLeadFrame (MLF) packages, September 2008. Online at www.amkor.com/index.cfm?objectid=42EDA4C7-5056-AA0A-E2A372F025BF8729, accessed 2013-10-28.

[CK08] I. Hakki Cavdar and Engin Karadeniz. Measurements of impedance and attenuation at CENELEC bands for power line communications systems. *Sensors*, 8(12):8027–8036, 2008.

[CUI13] CUI inc. VSK 15 W — AC-DC power supply, January 2013. Online at http://www.cui.com/product/resource/vsk-s15-series.pdf, accessed 2013-12-19.

[EN01a] EN. EN 50065-1. Transmission de signaux sur les résaux électriques basse-tension dans la bande fréquences de 3 kHz à 148,5 kHz — Partie 1: Règles générales, bandes de fréquences et perturbations électromagnétiques, April 2001.

[EN01b] EN. EN 50065-7. Signalling on low-voltage electrical installations in the frequency range 3 kHz to 148.5 kHz - Part 7: Equipment impedance, January 2001.

[ERD09] Électricité réseau distribution France, ERDF. Linky PLC profile functional specification, September 2009. Online at http://www.erdfdistribution.fr/medias/Linky/ERDF-CPT-Linky-SPEC-FONC-CPL.pdf, accessed 2014-02-20.

[Fut09] Future Technology Devices International (FTDI). An234 — FTDI drivers installation guide for Windows 7, October 2009. Online at http://www.ftdichip.com/Support/Documents/AppNotes/AN_119_FTDI_Drivers_Installation_Guide_for_Windows7.pdf, accessed 2013-11-15.

[Fut13a] Future Technology Devices International (FTDI). An234 — FTDI drivers installation guide for Windows 8, February 2013. Online at http://www.ftdichip.com/Support/Documents/AppNotes/AN_234_FTDI_Drivers_Installation_Guide_for_Windows_8.pdf, accessed 2013-11-15.

[Fut13b] Future Technology Devices International (FTDI). FTDI virtual COM port drivers, June 2013. Online at http://www.ftdichip.com/Drivers/VCP.htm, accessed 2013-10-10.

[IEC05] IEC. Effects of current on human beings and livestock — part 1: General aspects., July 2005.

[Mea12] Mean Well. PM-10 series — 10 W output switching power supply, November 2012. Online at http://www.meanwell.com/search/PM-10/PM-10-spec.pdf, accessed 2013-12-19.

[Mic13] Microsoft Corporation. Microsoft .NET framework 4.5 installer, September 2013. Online at http://www.microsoft.com/visualstudio/eng/downloads#d-net-45, accessed 2013-10-10.

[ON 02] ON Semiconductor. Board mounting notes for quad flat-pack no-lead package (QFN) (application note AND8086/D), August 2002.

[ON 11] ON Semiconductor. Thermal sensing methods used in on semiconductor devices (application note AND8494/D), October 2011. Online at http://www.onsemi.com/pub/Collateral/AND8494-D.PDF.

[ON 12] ON Semiconductor. NCN49599 evaluation kit (EVK) daughterboard (revision 2) schematics, December 2012.

[ON 13] ON Semiconductor. NCN49597 evaluation kit (EVK) daughterboard (revision 3) schematics, January 2013.

[ON 14a] ON Semiconductor. NCN49597 power line carrier modem datasheet, December 2014.

Online at http://www.onsemi.com/pub_link/Collateral/NCN49597-D.PDF.

[ON 14b] ON Semiconductor. NCN49599 power line carrier modem datasheet, December 2014. Online at http://www.onsemi.com/pub_link/Collateral/NCN49599-D.PDF.

[ON 14c] ON Semiconductor. PLC evaluation kit (EVK) motherboard (revision 6b) schematics, June 2014.

[Sto06] Roger Stout. *How To Use Thermal Data Found in Data Sheets (application note AND8220/D)*. ON Semiconductor, April 2006.

[Sto13] Roger Stout. Board mounting notes for WDFN8, 3.3 x 3.3 mm, 0.65P (8FL) (application note AND9138/D), May 2013.

[Sto14] Roger Stout. *Thermal Considerations for the NCS5651 (application note AND8402/D)*. ON Semiconductor, August 2014. Online at http://www.onsemi.com/pub_link/Collateral/AND8402-D.PDF.

[Tek05] Tektronix. Fundamentals of floating measurements and isolated input oscilloscopes, February 2005. Application note, online at http://info.tek.com/www-fundamentals-of-floating-measurements-and-isolated-input-oscilliscopes.html.

[VC11] Andrey Vangev and Peter Cox. *Transformers for PLC modems (application note)*. ON Semiconductor, March 2011.

[Ver15a] Michiel Verschueren. IEC Linky stack user manual v0.1 — draft. User manual, ON Semiconductor, Energy Management Software Engineering Team, 2015.

[Ver15b] Michiel Verschueren. ON PL110 stack user manual v0.1 — draft. User manual, ON Semiconductor, Energy Management Software Engineering Team, 2015.

[WVC14] Pierre Woestyn, Andrey Vangev, and Peter Cox. *Getting started with power line communication (application note AND9165/D)*. ON Semiconductor, June 2014. Online at http://www.onsemi.com/pub_link/Collateral/AND9165-D.PDF.

[WVC15] Pierre Woestyn, Andrey Vangev, and Peter Cox. *Mains synchronisation for PLC modems (application note)*. ON Semiconductor, August 2015.

Index