



Roy Gillotti <gillottir@gmail.com>

TWL6040 decoupling

15 messages

David Shah <dave@ds0.me>
To: hns@goldelico.com
Cc: kernel@pyra-handheld.com

Wed, Jul 29, 2020 at 10:42 AM

Hi Nikolaus,

While debugging the glitching audio noise issue (usually during loud volume playback), it is starting to seem like it is of electrical origin. I can trigger it consistently by enabling the vibrate while sound is playing (regardless of volume).

I also discovered that the TWL6040 has a test mode where it loops back downlink data to uplink.

I tested with this and the PDM data it was receiving was still correct, which pretty much eliminates a problem in the OMAP5 or between the OMAP5 and TWL.

This makes me think that high power consumption or electrical noise, due to loud volumes or the vibrate being turned on, is causing the TWL6040 to go haywire internally.

Comparing the Pyra schematics to the TWL6040 datasheet, it seems the 2.2uF capacitors on the HF amp and vibra power supplies are missing. Are you relying on decoupling elsewhere being enough? Do you think this could be the cause of the problem (I appreciate if it is then not much to be done about it now other than find the best possible software workarounds.)

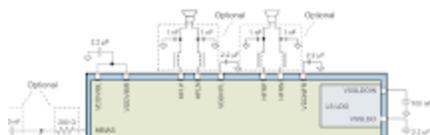
Best

David

Kernel mailing list

Kernel@pyra-handheld.com

<http://pyra-handheld.com/cgi-bin/mailman/listinfo/kernel>



twl6040_decoupling.png
39K

H. Nikolaus Schaller <hns@goldelico.com>

Wed, Jul 29, 2020 at 11:25 AM

To: David Shah <dave@ds0.me>

Cc: kernel@pyra-handheld.com

Hi David,

> Am 29.07.2020 um 16:42 schrieb David Shah <dave@ds0.me>:

>

> Hi Nikolaus,

>

> While debugging the glitching audio noise issue (usually during loud

> volume playback), it is starting to seem like it is of electrical

> origin. I can trigger it consistently by enabling the vibrate while

> sound is playing (regardless of volume).

I have tried and yes, I can trigger it sometimes through running the vibra through a second shell. Never did that before...

For me it happens only at higher volume levels. And sometimes it goes away if vibra stops. It always goes away if I stop aplay and restart.

So something gets "reinitialized".

Interesting observation!

>

> I also discovered that the TWL6040 has a test mode where it loops back > downlink data to uplink.

>

> I tested with this and the PDM data it was receiving was still correct,

> which pretty much eliminates a problem in the OMAP5 or between the

> OMAP5 and TWL.

Good to know.

>

> This makes me think that high power consumption or electrical noise,

> due to loud volumes or the vibrate being turned on, is causing the

> TWL6040 to go haywire internally.

Hm. Yes. The speakers can draw up to ca. 1 Ampere peak current from the battery. This gives a significant noise and ripple at higher volume...

>

> Comparing the Pyra schematics to the TWL6040 datasheet, it seems the

> 2.2uF capacitors on the HF amp and vibra power supplies are missing.

Indeed.

I have checked with the OMAP5 EVM and there, there are no capacitors as well.

Of course there are also no handsfree and speaker connector pins and the power supply is grounded :)

Next was to check the PandaBoard ES schematics.

It has a common 4.7uF for both channels. And only vibra power is

grounded.

So we probably did miss that when converting the OMAP5 EVM schematics into Pyra schematics and didn't compare with the PandaBoard and didn't look intensively into the twl6040 data sheets.

> Are you relying on decoupling elsewhere being enough?

A look into the PCB data shows that the nearest decoupling is a little far away. Mainly it goes through the board-to-board connector on the main board. Room for a lot of resistance and inductance. And electrical problems.

> Do you think this
> could be the cause of the problem (I appreciate if it is then not much
> to be done about it now other than find the best possible software
> workarounds.)

With this observation it is difficult to exclude this as the reason :)

If you see the loopback data being correct, I would suspect that something in the PWM system of the Class D driver gets out of sync. Maybe a clock loses some cycles or a (D)PLL becomes unstable. It must be something indirect because the scratchy noise usually (not always!) persists until you end the aplay command and start a new one.

What I have seen is that there is room for an optional LED on the backside of the CPU board directly connected to these pins. It was left there from initial hardware debugging. It should be not too difficult to try to add a 4.7uF or 10uF capacitor there to test the result.

I am not sure if that would be a good mass production solution, because it sticks to the battery bay plastics and/or the heat shield and slows down the production process and/or must be hand-soldered.

But there would be room and easy access on component side of the LC15 processor board to add such capacitor(s) by modified copper trace layout. So it could be something for the next mass production run (unless ED has already ordered thousands of PCBs).

Thanks for being such lucky to be the first observer of the relation between vibra and handsfree sound quality which seems to be the key to fix it.

>
> Best
>
> David
> <twl6040_decoupling.png>

Best regards,
Nikolaus

Kernel mailing list

Kernel@pyra-handheld.com

<http://pyra-handheld.com/cgi-bin/mailman/listinfo/kernel>

Michael Mrozek <EvilDragon@openpandora.org>

Wed, Jul 29, 2020 at 11:37 AM

To: "H. Nikolaus Schaller" <hns@godelico.com>, David Shah <dave@ds0.me>

Cc: kernel@pyra-handheld.com

Am Mittwoch, den 29.07.2020, 17:25 +0200 schrieb H. Nikolaus Schaller:

Hi,

> What I have seen is that there is room for an optional LED on the
> backside of the CPU board directly connected to these pins. It was
> left there from initial hardware debugging. It should be not too difficult
> to try to add a 4.7uF or 10uF capacitor there to test the result.

Oh, that's good news!

Can you let me know where to solder them so I'll try that here?

> I am not sure if that would be a good mass production solution, because it
> sticks to the battery bay plastics and/or the heat shield and slows down
> the production process and/or must be hand-soldered.

Well, we tested various materials between the CPU board and the plastics - but the best solution is to put NOTHING in between.

If we put a metal sheet in there, the heat goes into the metal and heats the battery - whereas almost nothing moves into the heatsink (where it should).

So there's a gap of 0.4mm between the CPU board and the backside, and as the keymat acts a bit like a cushion as well, I guess a small part like a capacitor could even use up to 0.6 to 0.7mm.

So there should be enough space to put a small capacitor in there.

> But there would be room and easy access on component side of the LC15
> processor board to add such capacitor(s) by modified copper trace layout.
> So it could be something for the next mass production run (unless ED has
> already ordered thousands of PCBs).

We have 500 already-produced PCBs - so I guess we'd need to fix these manually. But hey, MWeston had to replace 500 capacitors on the Pandora PCBs manually as well.

Only bigger downside would be that we have another setup costs for the next CPU boards... that's around 5000 EUR altogether :/

--

Greetings,

Michael Mrozek

OpenPandora GmbH
CEO: Michael Mrozek

Schäffbräustr. 11
85049 Ingolstadt
Germany
Tel.: +49 841 990 5548
<http://www.openpandora.de/>
HRB 4879, Amtsgericht Ingolstadt

eMail: mrozek@openpandora.org

[Quoted text hidden]

David Shah <dave@ds0.me>

Wed, Jul 29, 2020 at 12:04 PM

To: Michael Mrozek <EvilDragon@openpandora.org>, "H. Nikolaus Schaller" <hns@goldelico.com>

Cc: kernel@pyra-handheld.com

Hi all,

I tried adding a 470nF 0402 capacitor in what I hope was the right place (see attached), and replacing the associated resistor with a short. Sadly that made no difference, but maybe 470n isn't high enough (I also have 4.7u 0805 but I suspect that might be too tall.

I'm still proud to own the first modded Pyra, even though it didn't work :D

Best

David

[Quoted text hidden]

Kernel mailing list

Kernel@pyra-handheld.com

<http://pyra-handheld.com/cgi-bin/mailman/listinfo/kernel>



photo_2020-07-29_17-03-18.jpg

171K

David Shah <dave@ds0.me>

Wed, Jul 29, 2020 at 12:14 PM

To: Michael Mrozek <EvilDragon@openpandora.org>, "H. Nikolaus Schaller" <hns@goldelico.com>

Cc: kernel@pyra-handheld.com

Tried again with the 4.7u 0805, it did fit with the case but still no difference.

Best

David

[Quoted text hidden]

[Quoted text hidden]

H. Nikolaus Schaller <hns@goldelico.com>

Wed, Jul 29, 2020 at 12:58 PM

To: David Shah <dave@ds0.me>

Cc: kernel@pyra-handheld.com

> Am 29.07.2020 um 18:04 schrieb David Shah <dave@ds0.me>:

>

> Hi all,

>

> I tried adding a 470nF 0402 capacitor in what I hope was the right
> place (see attached),

yes, from a first glance it looks to be right!

> and replacing the associated resistor with a
> short. Sadly that made no difference, but maybe 470n isn't high enough
> (I also have 4.7u 0805 but I suspect that might be too tall.

Hm. I hope I find some minutes to open my sample and try the same.
Then we know.

>

> I'm still proud to own the first modded Pyra, even though it didn't
> work :D

++

BR, Nikolaus

[Quoted text hidden]

> <photo_2020-07-29_17-03-18.jpg>

[Quoted text hidden]

David Shah <dave@ds0.me>

Wed, Jul 29, 2020 at 1:15 PM

To: "H. Nikolaus Schaller" <hns@goldelico.com>

Cc: kernel@pyra-handheld.com

Playing around a bit more. I think there is a small improvement -
whereas before the glitching would continue 10-15s after vibra stopped,
now it has about a 70% chance of fixing itself once the vibra stopped.

It would also be good to see if the decoupling improves the glitching
when is volume too loud issue (without the vibra). Unfortunately, this
rarely happened anyway on my Pyra so I can't comment without a lot more

testing.

Best

David

[Quoted text hidden]

H. Nikolaus Schaller <hns@goldelico.com>

Wed, Jul 29, 2020 at 1:19 PM

To: David Shah <dave@ds0.me>

Cc: kernel@pyra-handheld.com

> Am 29.07.2020 um 18:58 schrieb H. Nikolaus Schaller <hns@goldelico.com>:

>

>

>> Am 29.07.2020 um 18:04 schrieb David Shah <dave@ds0.me>:

>>

>> Hi all,

>>

>> I tried adding a 470nF 0402 capacitor in what I hope was the right

>> place (see attached),

>

> yes, from a first glance it looks to be right!

>

>> and replacing the associated resistor with a

>> short. Sadly that made no difference, but maybe 470n isn't high enough

>> (I also have 4.7u 0805 but I suspect that might be too tall.

>

> Hm. I hope I find some minutes to open my sample and try the same.

> Then we know.

Here is a full picture where I tried to mimic the colours:

Maybe the copper trace between the C and the R is too small. So a wire directly across all three solder pads could be better. Or scratching off the solder mask from the ground plane around and soldering the capacitor directly to GND.

If I find time (which is unfortunately unlikely) I could add two wires to the pads of the capacitor and get them outside the Pyra to be able to measure the voltage with the oscilloscope.

Some thoughts about the software side: I remember that Peter Ujfalusi long time ago mentioned different clocks. Maybe the 32kHz clock or the McPDM clock which could be used as a reference. And there was a kernel version where sound became a little better (but I wouldn't bet for it). So there may still be another effect in software. Maybe in addition so that what we observe is the sum of several effects.

One question about the loopback: does it show if the PCM data is coming in the right and stable timing? E.g. speed may fluctuate but still provide perfect data mirroring. Nevertheless, fluctuating speed may become audible.

But it is difficult to see how this would be influenced by the Vibra. Well, the Vibra is also controlled through some PCM channel because the motor can "play" the bass sound...

>>

>> I'm still proud to own the first modded Pyra, even though it didn't

>> work :D

>

> ++

>

> BR, Nikolaus

>

>>

>> Best

>>

>> David

>>

>> On Wed, 2020-07-29 at 17:37 +0200, Michael Mrozek wrote:

>>> Am Mittwoch, den 29.07.2020, 17:25 +0200 schrieb H. Nikolaus

>>> Schaller:

>>>

>>> Hi,

>>>

>>>> What I have seen is that there is room for an optional LED on the

>>>> backside of the CPU board directly connected to these pins. It was

>>>> left there from initial hardware debugging. It should be not too

>>>> difficult

>>>> to try to add a 4.7uF or 10uF capacitor there to test the result.

>>>>

>>>> Oh, that's good news!

>>>> Can you let me know where to solder them so I'll try that here?

>>>>

>>>> I am not sure if that would be a good mass production solution,

>>>> because it

>>>> sticks to the battery bay plastics and/or the heat shield and slows

>>>> down

>>>> the production process and/or must be hand-soldered.

>>>>

>>>> Well, we tested various materials between the CPU board and the

>>>> plastics - but

>>>> the best solution is to put NOTHING in between.

>>>>

>>>> If we put a metal sheet in there, the heat goes into the metal and

>>>> heats the

>>>> battery - whereas almost nothing moves into the heatsink (where it

>>>> should).

>>>>

>>>> So there's a gap of 0.4mm between the CPU board and the backside, and

>>>> as the

>>>> keymat acts a bit like a cushion as well, I guess a small part like a

>>>> capacitor

>>>> could even use up to 0.6 to 0.7mm.

>>>>

>>>> So there should be enough space to put a small capacitor in there.

>>>>

>>>>> But there would be room and easy access on component side of the

>>>> LC15
>>>> processor board to add such capacitor(s) by modified copper trace
>>>> layout.
>>>> So it could be something for the next mass production run (unless
>>>> ED has
>>>> already ordered thousands of PCBs).
>>>>
>>>> We have 500 already-produced PCBs - so I guess we'd need to fix these
>>>> manually.
>>>> But hey, MWeston had to replace 500 capacitors on the Pandora PCBs
>>>> manually as
>>>> well.
>>>>
>>>> Only bigger downside would be that we have another setup costs for
>>>> the next CPU
>>>> boards... that's around 5000 EUR altogether :/
>>>>
>>>>
>> <photo_2020-07-29_17-03-18.jpg>
>
>
> _____
> Kernel mailing list
> Kernel@pyra-handheld.com
> <http://pyra-handheld.com/cgi-bin/mailman/listinfo/kernel>

Kernel mailing list
Kernel@pyra-handheld.com
<http://pyra-handheld.com/cgi-bin/mailman/listinfo/kernel>

 **Ohne Titel.png.pdf**
247K

David Shah <dave@ds0.me>
To: "H. Nikolaus Schaller" <hns@goldelico.com>
Cc: kernel@pyra-handheld.com

Wed, Jul 29, 2020 at 1:32 PM

Hi Nikolaus,

Couple of thoughts below.

[Quoted text hidden]

Yes, that matches what I did.

>
> Maybe the copper trace between the C and the R is too small. So a
> wire directly
> across all three solder pads could be better. Or scratching off the
> solder mask from
> the ground plane around and soldering the capacitor directly to GND.
>

In my second attempt with 4.7u, I did effectively end up with one solder blob all the way across that should be pretty low impedance.

- > If I find time (which is unfortunately unlikely) I could add two
- > wires to the pads of the
- > capacitor and get them outside the Pyra to be able to measure the
- > voltage with the oscilloscope.
- >
- > Some thoughts about the software side: I remember that Peter Ujfalusi
- > long time ago mentioned
- > different clocks. Maybe the 32kHz clock or the McPDM clock which
- > could be used as a
- > reference. And there was a kernel version where sound became a little
- > better (but I wouldn't
- > bet for it). So there may still be another effect in software. Maybe
- > in addition so that
- > what we observe is the sum of several effects.
- >

Interesting, the clocking definitely isn't quite right in the latest as high power mode also used to work on 4.19 but doesn't seem to any more.

- > One question about the loopback: does it show if the PCM data is
- > coming in the right
- > and stable timing? E.g. speed may fluctuate but still provide perfect
- > data mirroring.
- > Nevertheless, fluctuating speed may become audible.
- >

This is true, there is some possibility for "common mode failure" that would symmetrically affect both directions.

- > But it is difficult to see how this would be influenced by the Vibra.
- > Well, the Vibra
- > is also controlled through some PCM channel because the motor can
- > "play" the bass sound...
- >

The Vibra seems to be set to "Input FF" mode by default AIUI it is controlled by I2C not the PDM stream when using vibra.py. But you can set it to PDM mode in alsamixer which seems to turn it always on (presumably a 1 is sent in the PDM stream for that channel by default?) noise that always goes away as soon as you turn it off again rather than glitching the stream.

It also seems to differ significantly less depending on the volume setting - it isn't directly scaled by volume in the same way as Vibra. Right glitching is; as if it is affecting something later in the signal path.

Strange that an unconnected channel also makes a difference, but the H-bridge is still presumably still switching and any H-bridge does produce some shoot-through noise.

I am wondering now if some config issue is making the TWL much more sensitive than it should be for some reason.

Best

David

[Quoted text hidden]

David Shah <dave@ds0.me>
To: "H. Nikolaus Schaller" <hns@goldelico.com>
Cc: kernel@pyra-handheld.com

Thu, Jul 30, 2020 at 6:15 AM

Well, I seem to have found a solution. It is totally and utterly inexplicable, but there you go.

The OMAP5 TRM says that the input buffer for MCPDM_LB_CLK - the clock going out from the OMAP5 to the TWL, that clocks the PDM stream and mirrors the TWL clock output - must be enabled; as the input signal is used internally for retiming purposes.

But using devmem2 to DISABLE this input buffer, against the TRM advice, seems to fix the stream desynchronisation issue!

```
# devmem2 0x4A0029A0 w 0x0080118
```

There is still a quiet a crackle when the vibra turns on or off, which really does seem like switching noise from the motor, but the stream is never actually corrupted and the crackle does not disrupt what is playing!

I'm done with OMAP stuff for today :) but I'll experiment more over the weekend. If you get a chance, I would be interested to see if this is reproducible on your unit, too.

I need to think a bit more about how this could actually look from a silicon point of view to see if there is a better fix. It is possible that disabling the input buffer means that the OMAP5 is using some internal clock instead that isn't disrupted by noise on the pad.

It could also be related to OAMP5 errata i800 / TWL errata 6, which involves relative clock phases. I had ruled that out because it would be odd for power supply noise to affect that; and it reads more like dropped/duplicated samples than total loss of stream.

Best

David

[Quoted text hidden]

H. Nikolaus Schaller <hns@goldelico.com>
To: David Shah <dave@ds0.me>
Cc: kernel@pyra-handheld.com

Thu, Jul 30, 2020 at 6:20 AM

Wow!

Maybe you can ask Peter Ujfalusi as he has done a lot of previous work for twl audio.

I'll check all this as soon as I find time.

BR and thanks,
Nikolaus
[Quoted text hidden]

H. Nikolaus Schaller <hns@goldelico.com>
To: David Shah <dave@ds0.me>
Cc: kernel@pyra-handheld.com

Thu, Jul 30, 2020 at 6:37 AM

Yes, I can confirm the effect.

After doing devmem2 0x4A0029A0 w 0x0080118 sound output stops (but aplay continues to run without notice about buffer under/overflow or alike).

After restarting aplay sound is better. And vibra influence is no longer noticeable.

Restoring devmem2 0x4A0029A0 w 0x1080118 makes problem come back - even while aplay runs.

But still it depends on volume.

So high volume + 0x1080118 is unreliable.

Hm. Is there a clock polarity setting? Either in McPDM or the twl6040?

If clock polarity would be wrong, then it may depend on power supply noise if the data line is sampled correctly and the previous bit value creeps in or some random bit is read.

BR,
Nikolaus
[Quoted text hidden]

aTc <atc@k-n-p.org>
To: kernel@pyra-handheld.com

Thu, Jul 30, 2020 at 9:28 AM

I can confirm that it works, but only for the speaker output, the headphone output goes silent at the point where it would have normally glitched into white noise.

On 7/30/20 12:37 PM, H. Nikolaus Schaller wrote:

Yes, I can confirm the effect.

After doing devmem2 0x4A0029A0 w 0x0080118 sound output stops (but aplay continues to run without notice about buffer under/overflow or alike).

After restarting aplay sound is better. And vibra influence is no longer noticeable.

Restoring devmem2 0x4A0029A0 w 0x1080118 makes problem come back - even while aplay runs.

But still it depends on volume.

So high volume + 0x1080118 is unreliable.

Hm. Is there a clock polarity setting? Either in McPDM or the twl6040?

If clock polarity would be wrong, then it may depend on power supply noise if the data line is sampled correctly and the previous bit value creeps in or some random bit is read.

BR,
Nikolaus
[Quoted text hidden]

aTc <atc@k-n-p.org>
To: kernel@pyra-handheld.com

Thu, Jul 30, 2020 at 9:43 AM

Great, I can't replicate it after a reboot.

Some of the weirdness on my unit might possibly be because the right speaker doesn't make proper contact with the pads (I had to bend the pins a little to get it to output sounds at all)

On 7/30/20 3:28 PM, aTc wrote:

I can confirm that it works, but only for the speaker output, the headphone output goes silent at the point where it would have normally glitched into white noise.

[Quoted text hidden]

Gražvydas Ignotas <notasas@gmail.com>
To: "H. Nikolaus Schaller" <hns@goldelico.com>
Cc: kernel@pyra-handheld.com

Thu, Jul 30, 2020 at 10:06 AM

On Thu, Jul 30, 2020 at 1:37 PM H. Nikolaus Schaller <hns@goldelico.com> wrote:

Hm. Is there a clock polarity setting? Either in McPDM or the twl6040?
If clock polarity would be wrong, then it may depend on power supply noise if the data line is sampled correctly and the previous bit value creeps in or some random bit is read.

Some trivia: we had this exact same issue just after pandora's mass production start back in 2010...

<https://git.openpandora.org/cgi-bin/gitweb.cgi?p=pandora-kernel.git;a=commitdiff;h=d0b55dac46924d029d3e123adde5e82aba7613ed>

Gražvydas

Kernel mailing list
Kernel@pyra-handheld.com
<http://pyra-handheld.com/cgi-bin/mailman/listinfo/kernel>