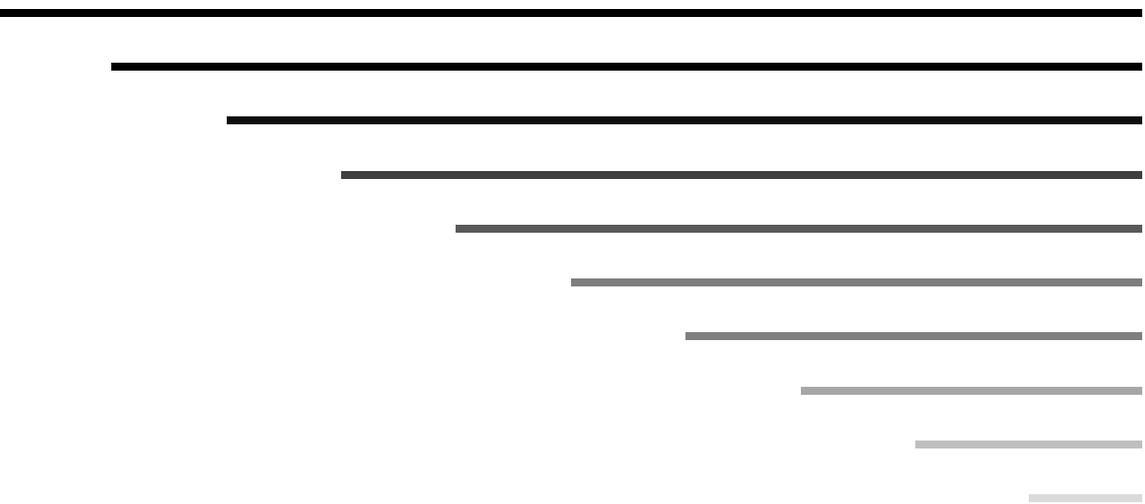




## μTasker Document

- **NXP LPC2XXX Tutorial – Ethernet and the Simulator**



## Table of Contents

1. Introduction.....	3
2. Getting Started .....	5
3. Testing on the Target.....	14
4. Compiling the Project for the Target.....	14
4.1. Compiling with IAR .....	15
4.2. Compiling with Rowley Crossworks .....	15
4.3. Compiling with GCC makefile.....	16
5. Downloading the Code to the Target.....	17
6. Watchdog and Debugging .....	19
7. Specific Notes concerning the LPC2XXX .....	20
8. A Tour through the Demo Web Server .....	21
8.1. Menu .....	21
8.2. LAN configuration .....	23
8.3. Serial Configuration .....	25
8.4. Statistics .....	26
8.5. I/O Page.....	27
8.6. Administration Side.....	27
8.7. LCD Side .....	29
8.8. Email Configuration Side .....	30
8.9. Dynamic Content Generation Side .....	31
9. Developing – Testing - Debugging .....	32
10. Off-Line Simulation.....	39
11. Conclusion.....	41

## 1. Introduction

μTasker is an operating system designed especially for embedded applications where a tight control over resources is desired along with a high level of user comfort to produce efficient and modular code.

The operating system is integrated with TCP/IP stack and important embedded Internet services alongside device drivers and system-specific project resources.

μTasker and its environment are essentially not hardware specific and can thus be moved between processor platforms with great ease and efficiency.

However the μTasker project setups are very hardware specific since they offer an optimal pre-defined (or a choice of pre-defined) configurations, taking it out of the league of “board support packages (BSP)” to a complete “project support package (PSP)”, a feature enabling projects to be greatly accelerated.

This tutorial will give you a flying start to using the μTasker on the NXP LPC2XXX. It can be fully simulated using VisualStudio 6.0 or higher and can be run on the various common evaluation boards (specific configurations for the Keil MCB2300 and Olimex LPC2378-STK evaluation boards are discussed in more detail), for which Rowley CrossWorks, IAR, Keil uVision and GCC project configurations are available. Building the project using GCC is included as an optional post-build step when working in VisualStudio, but debugging with the GNU tool chain is not discussed. Generally a JTAG debugger should be used together with the chosen development environment.

Although there are various ways to work with the LPC2XXX, the following methods are described in this tutorial:

- Programming of the chip is performed using FLASH MAGIC from Embedded Systems Academy (sponsored by NXP) via serial interface
- Debugging is performed using IAR Embedded Workbench and the JLink JTAG debugger or using Rowley CrossWorks and GNU compiler with CrossConnect JTAG debugger

The LPC2XXX range supported by the μTasker project stretches from the smallest LPC2101 (with 8k FLASH and 2k SRAM) up to the LPC24XX (with integrated LCD controller and external SDRAM interface). Based on the ARM7 core, the powerful LPC23XX is manufactured by NXP (formerly Philips) and was release in its first version around autumn 2006. The first silicon had a few problems which caused various difficulties for the first programmers trying to get to grips with the chips but the devices has matured in the meantime and are now readily obtainable in a number of derivatives. There are a few bugs for the unwary user, especially if the evaluation board is not brand new, and these will be pointed out later in the tutorial. The standard workarounds have already been prepared in the μTasker and these can be individually deactivated if it is clear that they are no longer an issue in latest devices. It is generally advisable to familiarise yourself with the latest errata from NXP and to carefully adhere to the guidelines in this tutorial: this should help avoid losing unnecessary time.

If you don't have the mentioned tools, other methods certainly are also possible, so use the tutorial as a base line and experiment with whatever tools you do own to produce equivalent results.

Some commercial compilers and debuggers are available as evaluation versions. Often they have a 32k code limit, which is suitable for the μTasker demo project – although it is not

possible to activate absolutely every feature at the same time, useful and interesting projects can be performed to see it in action.

Other commercial compilers only allow a smaller code size for evaluations. These are not mentioned anymore since 16k is not adequate to produce serious demonstrations of embedded IP work.

Even if you don't have the tools available at the moment, the μTasker simulator will enable you to get started. Therefore there should be nothing standing in your way to getting your first, powerful embedded IP project up and running.

This version of the SAM7X tutorial has been adapted for the μTasker V1.4 version including new graphical LCD support, which is also illustrated here. Some LPC24XX devices include an LCD controller and, when these types are selected (LPC2478FET208 and LPC2478FBD208) the demo project will also allow images for the LCD to be posted via web browser – see the LCD User's Guide

<http://www.utasker.com/docs/uTasker/uTaskerLCD.PDF> for more details about this capability.

For more μTasker documentation visit: <http://www.utasker.com/docs/documentation.html>

Note further that MODBUS support (serial and TCP) is available as an extension pack for the LPC2XXX – see the MODBUS user's guide for full details:

[http://www.utasker.com/docs/MODBUS/uTasker\\_MODBUS.PDF](http://www.utasker.com/docs/MODBUS/uTasker_MODBUS.PDF)

## 2. Getting Started

You are probably itching to see something in action and so why hang around. Let's start with something that will already impress you and your friends – no simple and basically useless demo which blinks an LED in a forever loop but something seriously professional and for real life projects very handy.

First I will assume that you have VisualStudio installed on your PC since we will first simulate everything – but don't worry, we are not going to see something attempting to interpret the instructions of the processor and requiring 2 minutes to simulate a couple of seconds of the application, instead we will see your PC operating in "real time" as the target processor. Your PC will not realise that the processor is simulated and so when you try contacting it by pinging it or browsing to it with your web browser, we will see that your PC will send IP frames to the network and will see answers on the network from the simulated device. Other PCs or IP enabled embedded devices on the network (as well as further simulated devices) will also be able to communicate with this simulated device. You can also capture the frames using a sniffer tool (we will use Wireshark) for further analysis and playback.... Getting excited? In a few minutes you will see it in action!

If you haven't VisualStudio (6.0 or newer) then there are trial versions and the free Express version [<http://msdn.microsoft.com/vstudio/products/trial/> or <http://www.microsoft.com/express/>] and you can probably pick up a 6.0 version for very little cash on Ebay. VS 6.0 is adequate for our work as are the Express editions. It contains a world class C-compiler and editor as well as loads of other tools which make it a must, even for embedded work.

The simulator requires also WinPCap to be installed (from <http://www.winpcap.org>) which is an industry-standard tool for link-layer network access in Windows environments. It is also used by the network sniffer Wireshark, which every Ethernet designer really should have. Therefore it is simplest to install Wireshark since this includes WinPCap, saving a step, and making for simplest installation. Wireshark can be downloaded from <http://www.wireshark.org/> and is discussed towards the end of the document.

If you don't want to see the simulator in action – **which would be a big mistake as you will miss the opportunity to save many hours of your own project time later** – there is target code which can be loaded to the target and will also run. This is also detailed later on in the tutorial.

You will find that the simulator is useful for most real embedded work. First, it allows testing things which you may not already have available as hardware (even if you have no evaluation board and cross compiler for it yet, you can start writing and testing your code) – it will allow you to use a matrix keyboard, LCDs, I<sup>2</sup>C EEPROM or SPI FLASH (and more) connected to the virtual target without having to get your soldering iron out to connect it. And it is so accurate that you can then cross compile to the target and it will (almost certainly) work there as well. This is the last time I will say it ... **don't make the BIG MISTAKE of taking a short cut when starting and diving into coding on the target.** If you are an embedded SW professional you will be missing the chance of saving enough time in a year for a couple of months extra vacation. If you are a hobby user, you will be missing the change to get out more...!

**So...**

**...fasten your seat belts since we are about to roll**

1. Simply copy the complete µTasker folder to your PC and go to `\Applications\uTaskerV1.4`. This is a ready to run project directory showing a useful application using network resources (and more).
2. Move to the sub director `Simulator` and open the VisualStudio project workspace `uTaskerV1-4.dsw`. This project is in the VisualStudio 6.0 format to ensure compatibility and, if you have a higher version, simply say that it should be converted to the new format – no problems are involved and it should build with no warnings.
3. *Ensure that the compiler is set up for the LPC2XXX target in the project's pre-compiler settings: look for the pre-processor define **\_LPC23XX**. If instead you find that the project is set up for another target, eg. `_HW_NE64` or `_M5223X` simply overwrite this with **\_LPC23XX**.*

Ensure that the project is configured for your target board in `config.h`: enable the chip type `LPC2378FBD144` if you are using the MCB2000 or Olimex LPC2378-STK. Select the board type too: `KEIL_MCB2300` or `OLIMEX_LPC2378_STK`. *Further boards can be added later if necessary to suit your own hardware.*

Build the project (use F7 as short cut) and you should find that everything compiles and links without any warnings.

Note that the VS project as delivered has two configurations (`uTasker - Win32 debug` and `uTasker - Win32 uTasker LPC23XX plus GNU build`). When the configuration with GNU build is selected a post build step will cross compile the project code using the GCC compiler. This is only possible when the GNU compiler is installed on the local PC and it may also be necessary to modify the paths in the bat file `\Applications\uTaskerV1.4\GNU_LPC23XX\Build_LPC23XX.bat` to suit. This is discussed in more detail in chapter 4.3.

4. I would love to be able to say “Execute” but there are a couple of things which have to be checked before we can do this. First of all, you will need to be connected to a network, meaning that your PC must have a LAN cable inserted and the LAN must be operational – either connected to another PC using a crossover cable or to a router or hub (wireless links tend not to work well but this may improve as the WinPCap capabilities get better with newer version).

Then we have to be sure that the network settings allow your PC to speak with the simulated device. The IP address must be within the local subnet:

Open the C-file `application.c` in the `uTaskerV1.4` project directory and check that the following default settings match your network settings (check what your PC uses in a DOS window with “ipconfig”)

```
static const NETWORK_PARAMETERS network_default = {
(AUTO_NEGOTIATE | RX_FLOW_CONTROL), // usNetworkOptions - see
                                     driver.h for other possibilities
{0x00, 0x00, 0x00, 0x00, 0x00, 0x00}, // ucOurMAC - our default MAC
{ 192, 168, 0, 3 }, // ucOurIP - our default IP address
{ 255, 255, 254, 0 }, // ucNetMask - our default network mask
{ 192, 168, 0, 1 } // ucDefGW - our default gateway
};
```

Just make sure that the network mask matches, that the device's IP address is within the local network and that the IP address defined doesn't collide with another one on the local network.

After making any modifications, simply compile the changes and we will be in business.

5. So now I can say EXECUTE (use F5 as short cut).

You will see the simulated device working away on your PC screen. If the define `SUPPORT_GLCD` hasn't been removed in `config.h` there will also be a simulated graphic LCD displaying some images. *This should be removed later for testing on the target if there is no graphical LCD actually connected.* There will also be one port output toggling away. Let's take a quick look at it. This is in fact the watchdog routine which is called every 200ms which is also toggling the output so that we can see that all is well (P2.0).

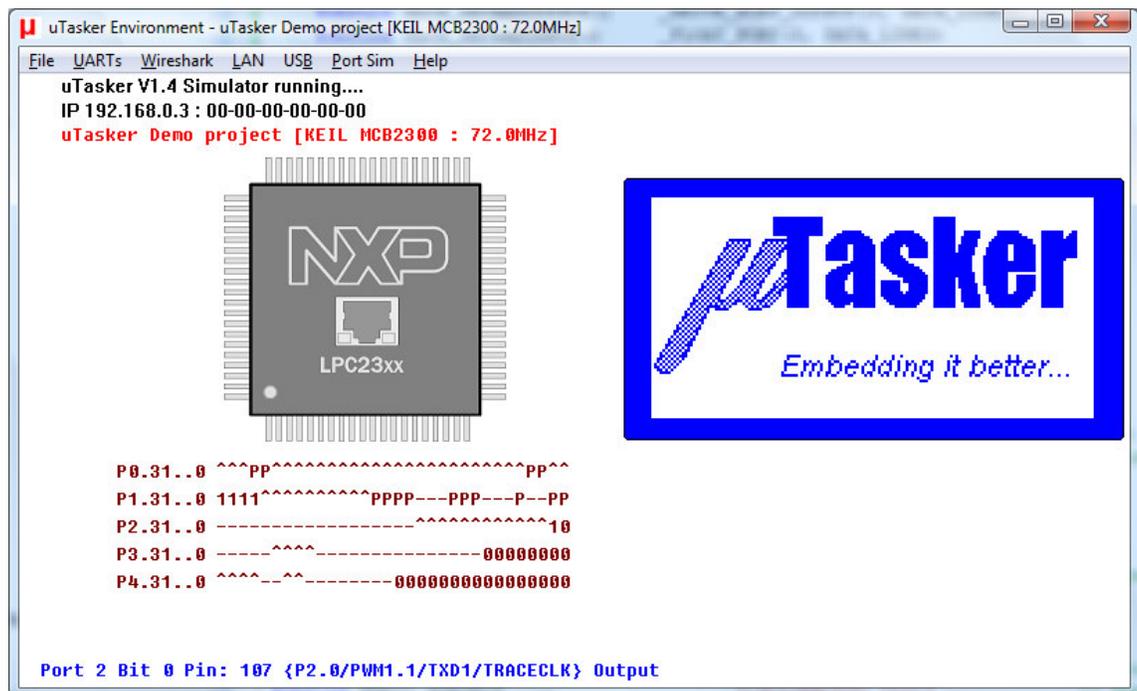


Figure 1. μTasker simulator running

If you hover your mouse over the port which is toggling, or other available ones, you will see which pin number it has on the device, its name and present use as well as the possible peripheral functions which it can be programmed to perform. In the screen above this is seen by the line in blue at the bottom of the simulation window.

Open the file `uTasker\Watchdog.c` (in the VisualStudio project manager). Put your cursor on the call to retrigger the watchdog "`fnRetriggerWatchdog()`" and hit the F9 key (*set a break point*). Almost immediately the program will halt and you can use F11 to step into the hardware specific routine responsible for triggering the watchdog and in our case also toggling the port output. The routine writes to the corresponding device registers and returns. Remove the breakpoint and let the simulated device run again using F5.

6. OK. So at least you are convinced that it is really running our project code, but it is not exactly something to write home about.

Let's get down to more serious stuff:

There are some menu items in the μTasker environment simulation window. Open "LAN | Select working NIC" and select the network card you would like the simulator to use – don't worry, it will be shared with anything else which is already using it.

Open a DOS window and do the standard PING test - "ping 192.168.0.3" if you didn't need to change anything, or the IP address you defined if you did.

The simulated device should now be responding. That means that your PC has sent PING test messages to the network and received answers from some network device (or course our internal simulated device – but no one will know the difference).

You can also try sending the ping test from another PC on the network and it will also receive an answer.

If you have a network sniffer you can also record the data from the network – don't worry if you haven't used a network sniffer before because we will come back to this later on and use Wireshark to do the job.

*Watch the LAN indicators on the processor's virtual LAN connector. The left one will blink green when a frame is received – if it is accepted by the device (matching address, broadcast or all when in promiscuous mode). The right one will blink red whenever the simulated device sends a frame to the network.*

7. I suggest that you now close the μTasker environment simulation window using the 'normal' method File | Exit or by clicking on the close cross in the top right hand corner.

This standard termination will cause the selected NIC to be saved to a file in the project simulation directory called **NIC.ini** and so you will not have to configure it the next time you start the μTasker environment simulation.

8. Now I'm sure you are not yet satisfied with the progress, so let's execute the project again (short cut F5) and this time we will do something a bit more interesting. Start your web browser (Internet Explorer, Firefox, Chrome, etc.) and establish a connection with the simulated target: In the URL line type in <http://192.168.0.3> (assuming its address has not been changed) and hit the ENTER key. The web browser will request the start page on the web server, which will return the following page the first time used:

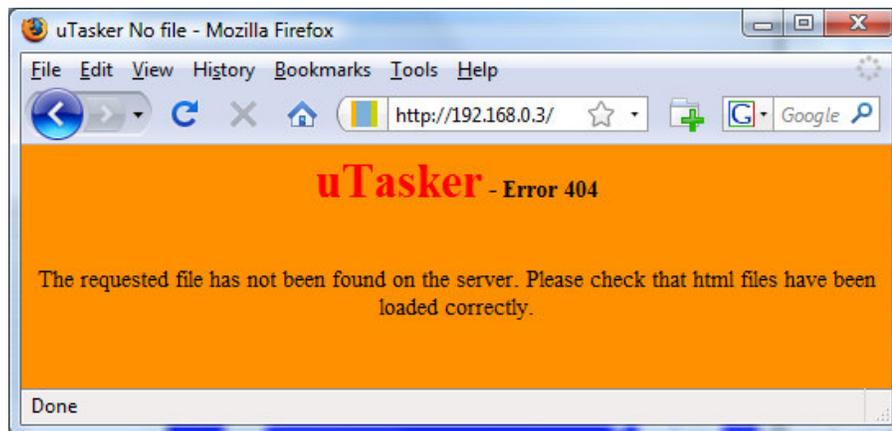


Figure 2. μTasker web server displaying that it has no web pages loaded

This is displaying that the page has not been found on the server, which is simply due to the fact that we haven't loaded any web pages to it yet, so we will do this as next step.

Note however that the web browser is displaying the NXP favicon. This is because this favicon is embedded in the code and not in the μFileSystem – this is explained in detail in the user file document

<http://www.utasker.com/docs/uTasker/uTaskerUserFiles.PDF>

9. Take a look in the project directory

\Applications\utaskerV1.4\WebPages\WebPagesLPC23xx\WebPagesDemo and you will find the web pages prepared for the demo project. These are HTM, JPG and GIF files. It is possible to load these pages using FTP clients but these tend to work quite differently from another and require various setups before they work smoothly. See the following document for more information on this subject:

<http://www.utasker.com/docs/uTasker/uTaskerFTP.PDF>

This is the reason that working with simple DOS FTP tends to be the fastest and simplest method and the bat file `Copy_all.bat` even does all the work for you. Check that its content is set to use the correct IP address and then double-click the bat-file to start the transfer of all of the web pages to the simulated target device. The transfer takes just a few seconds.

Then go back to the web browser window and make a refresh. Now you will see the start side and can navigate to a number of other pages.

Before we get into any more details about the web pages and their uses, please modify and save the Device ID (on the start page – default “uTasker Number 1”) to any name of your choice and close the μTasker environment simulation window by using the ‘normal’ exit method.

By doing this, the contents of the device's simulated FLASH (Flash is used to save the web pages in the μFileSystem and also any parameters, for example the device ID which you have modified) are saved to a file in the project simulation directory called **FLASH\_LPC23XX.ini**.

Now execute again (F5) and check that the loaded web pages are all there and that after a refresh of the start side also the device ID which you chose is correct.

Now you should understand the operation of the simulator; *on a normal exit it saves all present values and settings and on the next start the device has the saved FLASH contents as at the last program exit*. This is exactly how the real device works since, after a reset, its FLASH contents are as saved – I mean, that is what FLASH is all about.

If however you have modified FLASH contents and do not want them to be saved, you can always avoid this by quitting the debugger (short cut SHIFT F5) which causes the simulator to be stopped without performing the normal save process. If you wish to start with a fresh device (with blank FLASH contents) then simply delete the **FLASH\_LPC23XX.ini** file from the simulator directory...

*Please review the notes in chapter 7 concerning the large granularity file system as used per default in LPC23XX projects before commencing your own work using web pages!*

10. There is one important setting which we should perform before continuing; that is to set a MAC address to our new device.  
So browse to the “LAN configuration side”, where you will see that the default MAC address is 00-00-00-00-00-00. This works, but we really should set a new one – as long as we are not directly connected to the Internet any non-used value can be set. You make one up. For example set 00-11-22-33-44-55 (watch that the entry format is correct and always start with 00!) and then click on “modify / validate settings”. The entry field will be set inactive since it is no longer zero and cannot be changed a second time (should you want to reset the value before we commit it to (simulated) FLASH click on “Reset changes”, otherwise click on “Save changes”).

The device is commanded to be reset at this point after a short delay of about 2 seconds to allow it to complete the web page serving. The simulator will open a dialog box to indicate that it has been terminated with a reset:

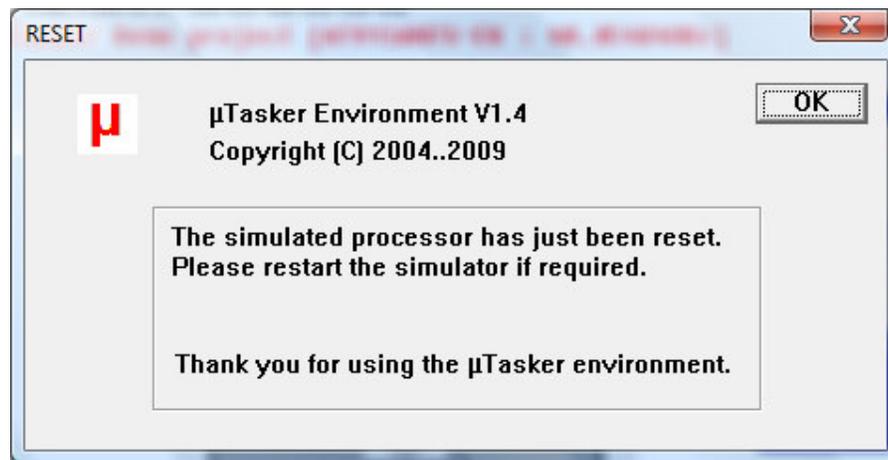


Figure 3. μTasker simulator indicating that the processor has just been reset

You can restart the simulator after the reset (which was interpreted as a ‘normal’ program exit and so the changes to the FLASH have also be saved to the hard disk and will be restored on the next program start) by using F5 as normal.

11. You may have read on the LAN configuration side that the new setting should be validated within 3 minutes – this is a safety mechanism so that falsely set critical values do not leave a remote device unreachable and means that we should establish a connection within three minutes of the new start to verify that all is well, otherwise the device will automatically delete the newly set values, reload the original ones and can then be contacted as before, using the original settings. So we will now validate our new setting (new MAC) so that it will always be used in the future.

Refresh the web pages by clicking first on “Go back to menu page” and then “Configure LAN Interface” again.

*You may well have a shock because it doesn't respond...but don't get worried because your PC may still be trying to reach the IP address using the previous MAC address, which is no longer valid (hence an incorrectly configured device can become unreachable). In a DOS windows type in “arp -d” which will delete the PC's ARP table – which is mapping IP address to MAC addresses and then it will work on the second attempt. Vista PCs may not always accept this command so wait a couple of minutes before starting the simulator again, until its ARP table automatically times out.*

The page will not allow any parameters to be modified and also the check box “*Settings validated*” shows that the device is waiting for the new values to be validated. To perform this, click on “*Modify / validate settings*”, after which the parameters will be displayed as validated. (*Don't forget to terminate the simulator normally later so that it is really the case*).

Should the 3 minutes have elapsed before you managed all that, the simulated device will again reset, clearing the temporary settings and reverting back to the initial MAC address. Just start the simulator again and repeat this step.

*While the project is waiting for the critical network changes to be validated you will notice that the LCD doesn't display anything. As soon as the validation sequence has been completed also the LCD display demonstration will start running as normal.*

12. Now you may be thinking that it is all well and good having a web server running on a device, sitting somewhere on the Internet but you probably want to use the device for controlling something - in its simplest form by switching something on or off. Using a relay this could be something quite powerful which can also have really useful benefits; imagine browsing to your device and commanding it to open the blinds or turn on the lawn sprinkler...the list of possibilities is almost endless... Look at the simulated LPC23XX and imagine that you have connected your lawn sprinkler to the output Port P2.3. Presently it is at the state input '1' [signified by '^'], the default state when the software starts since there is a pull-up in the device. Now click on the link menu page to “*Open an I/O window*” – it is shown in the following screen shot. Notice that the port P2.3 is being displayed as an input and its digital value as active ('1') in its check box.

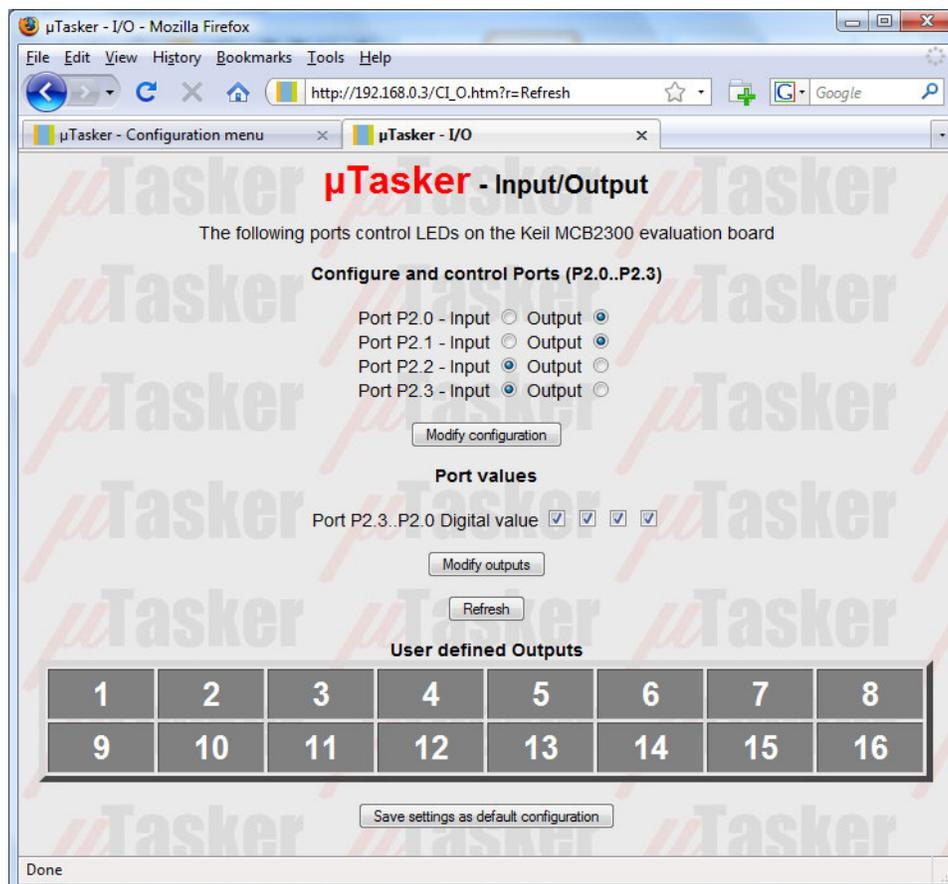


Figure 4. μTasker web server showing the I/O control page

Do a quick test: Hover your mouse over the port P2.3 in the simulator – it will tell you which pin it is on your device. Now click – this has toggled the simulated input (each time you click, it will change state again) – leave it at the low state ‘0’, displayed as ‘v’. Now refresh the I/O web page by clicking in the Refresh button. What is the web page now displaying? If it is behaving correctly it will of course be displaying the new port input state.

Now we can configure it as an output and control its state from the I/O page:

Set P2.3 to be an output, rather than an input and click on “*Modify configuration*”. Afterwards set the port state to ‘1’ by setting the first port value from the left and click on “*Modify outputs*”. Watch the state of Port P2.3 on the simulated device – you will see that the output is now set and now your lawn would get the water it has been waiting for!

Set the state to ‘0’ and it will be turned off. In this example web page, all of the 4 Ports P2.3..P2.0 outputs can be controlled individually.

P2.0 is used as a RUN LED in the demo project and so is configured as an output by default, as is P2.1.

If you would like a certain setting to be returned after every reset of the device, simply save the setting by clicking on “*Save settings as default configuration*” and you will see that they are indeed set automatically when the device is started the next time.

The demo project includes also 16 user controllable outputs. These are automatically configured as outputs in the project code and can be seen as outputs with state ‘0’ when looking at the simulator. If you now click on the user defined outputs on the I/O page you will see that they immediately toggle their port output and the state is either displayed as a grey button when ‘0’ or as red buttons when ‘1’. Also the present state of the 16 outputs can be saved as default states when the target starts the next time. *Note that these outputs can be mapped to any port and port bit in the demo project – it shows a useful example of flexible independent output control.*

13. The last setting which we will change before taking a look at the project on the target hardware is to activate HTTP user authentication and disable the FTP server so that no one else without our password has entry to the web server and no sneaky person can change the web pages which we have just programmed...

Browse to the administration page, deactivate FTP and activate HTTP server authentication before selecting the action to “*Modify and save server settings*”. Finally click on “*Perform desired action*” and close your web browser..

Open the web browser again and enter the device’s address <http://192.168.0.3> (or your address) and this time you will need to enter the user name “ADMIN” and password “uTasker” to get in.

Try to do something with the FTP server (like transfer the file with *Copy\_all.bat* or delete the file system contents with *delete\_all.bat*) and you will find that it simply won’t work anymore. This is because the FTP server is no longer running in the simulated device and you can rest assured that no one out there will be able to change anything which you have loaded.

*You probably noticed a button at the bottom of the administration page which allows a file to be searched for and uploaded to the target as a new software version. This is an example of using the HTTP POST method to transfer data to the embedded*

device. The μTasker project for the LPC2XXX supports the upload of complete new software using this method (as well as by FTP). For further details about using this capability on the LPC2XXX see the document [http://www.utasker.com/docs/uTasker/uTaskerBoot\\_003.PDF](http://www.utasker.com/docs/uTasker/uTaskerBoot_003.PDF) :You will see that you will then be able to develop code for a real target on the internet, including uploading it from anywhere where you have web access to it!

The uTasker V1.4 includes graphical LCD support, which has been seen above. The MCB2300 board has a character LCD, which can be simply configured in `config.h` by activating `SUPPORT_LCD` and disabling `SUPPORT_GLCD`. In this case the simulator will show the character LCD.

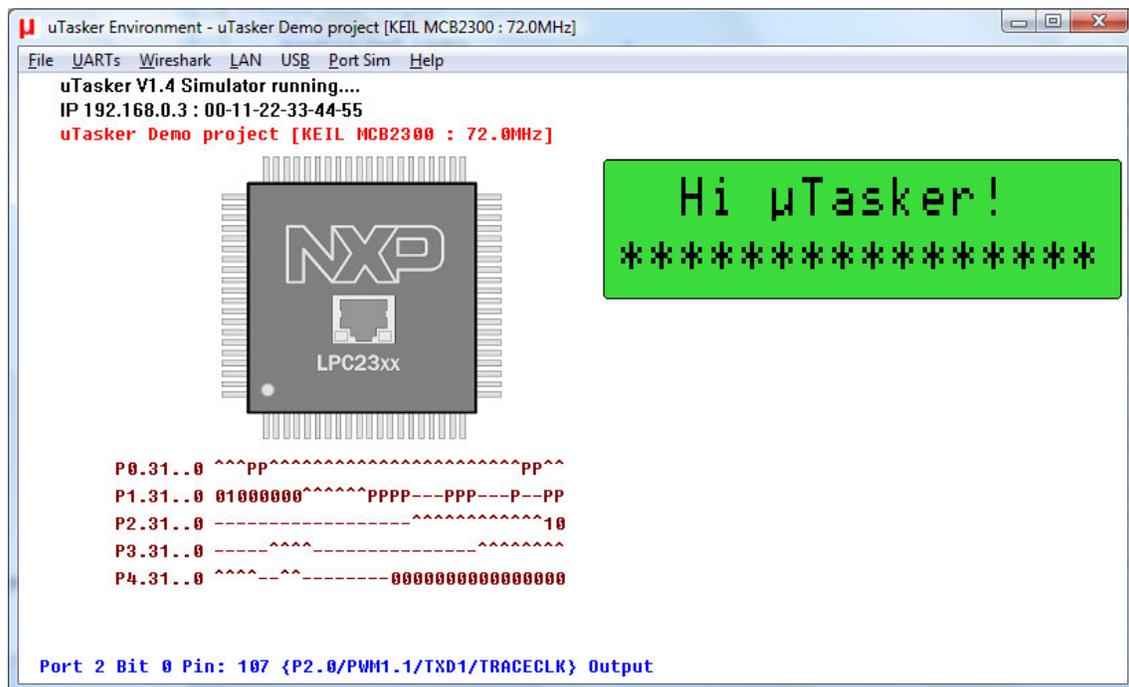


Figure 5. μTasker simulator showing alternative character LCD as available on the MCB2300 evaluation board

I hope that the introduction has shown that we are dealing with something which is very simple to use but is also very powerful in features. There are lots of details which need to be learned to understand and use everything to its fullest capabilities and it is up to you to decide whether you want this or would like to simply use the μTasker and its capabilities as a platform for your own application development. In any case you have just learned the basics of a powerful tool which is not only great fun but can really save you lots of development time.

### 3. Testing on the Target

Up until now we have been using the simulator and hopefully you will agree too that it is a very useful tool. We have tested some quite useful code designed to run on the LPC23XX and done this in an environment enabling us to do embedded IP stuff in real time. We haven't actually done any debugging or added new code but you can probably imagine the advantages of being able to write and test it on the PC before moving to the target.

At the same time you are probably thinking about how that will all work on the real device. Perhaps you are worried that it is all a show after all and the real device will remain as dead as a door nail or at the best crash every time the user breaths heavily. So let's prove that this is not the case by repeating the first part of the tutorial, but this time we will go live...

#### Notes on Ethernet PHY

We will use the Keil MCB2300 and/or Olimex LPC2378-STK evaluation board(s). The MCB2300 uses the National Semiconductor physical layer Ethernet device and the Olimex uses the Micrel KS8721 – if you have a different board using these parts the code should also run on it. If you have a different PHY you may have to adjust the PHY address and ID to get it recognised – please ask in the μTasker forum if you have a specific requirement.

### 4. Compiling the Project for the Target

**Warning:** the μTasker demo project V1.4 is delivered with graphical LCD support activated. This can cause the SW to hang if there is no real LCD connected. Before compiling for the target it is recommended to disable `SUPPORT_GLCD` in `config.h` so that this cannot occur.

The first thing that we need to do is to compile the project for the target. This means that we will be cross compiling the code which we already have been testing with the simulator, using an assembler, a C-compiler and a linker. There are a number of such tools available and unfortunately they are not all compatible in every aspect, even if ANSI compatible. Basically there is always the chore of getting the thing to reset from the reset vector, meaning that the start-up code must be available and at the correct location. Then there are specialities concerning how we force certain code or variables to certain addresses or regions and how we need to define interrupts routines so that they are handled correctly. Sometimes there is also need for some special assembler code which does such things as setting the stack pointer or enabling and disabling interrupts.

The μTasker demo project for the LPC2XXX is delivered with projects for IAR Embedded Workbench project, Rowley Associates Crossworks and Keil uVision as well as a 'stand-alone' GNU build. The latter is integrated in the VisualStudio project as an optional post-build step. This requires a GNU compiler to be available on the local PC – it will compile the project but doesn't support debugging.

The IAR compiler is a very good compiler, producing highly efficient code so is recommended for professional work. There are evaluation versions available from IAR - which are generally time or code size limited - for anyone wanting to get a feel for the product. <http://www.iar.com/>

The Rowley Associates CrossWorks is GNU compiler based but supplies a complete development environment including support. This is recommended for professional users who would like to base their work on GNU but don't want the hassle of having to configure the various tools before being able to start productive work. They also offer attractively priced educational and personal licenses. <http://www.rowley.co.uk/>

Otherwise it should be possible to get the code up and running with any other compiler but it will involve setting up an equivalent project and a certain amount of knowledge of the processor being used and various details about the compiler itself.

By the way, if you can't wait until you have installed the compiler or get your own project up and running, there are files delivered with the μTasker which were compiled with the GNU compiler and can be loaded already:

### 4.1. Compiling with IAR

You can find the project and the target file(s) in the sub-directory called `IAR5_LPC23XX`:

`uTaskerV1.4.eww` This is the IAR Embedded Workbench project work space

`IAR5_LPC23XX\Release\Exe\uTaskerV1.4.hex` This is the file which you can download to the target using FLASH Magic, as explained later

*When compiling the project, expect three non-serious compiler warnings but no errors.*

### 4.2. Compiling with Rowley Crossworks

You can find the project and the target file(s) in the sub-directory called `Rowley_LPC23XX`:

`uTaskerV14.hzp` This is the Rowley CrossStudio project work space

`Rowley_LM3SXXXX \THUMB Flash Release\Exe\uTaskerV14.bin` This is the file which you can download to the target using FLASH Magic, as explained later

*When compiling the project, expect neither warnings nor errors to be emitted by the GNU compiler.*

### 4.3. Compiling with GCC makefile

The µTasker simulator runs under VisualStudio and has two build configurations. The default is **Debug**, which we have been using up to now. The second is **uTasker LPC23XX with GNU build**. This does the same as the first configuration but automatically starts a GNU post-build when the VisualStudio project builds successfully. This is simply the call of a bat file which starts the GNU build, but the advantage is that the output of the GNU build appears in the compile windows.

To activate the GNU build option you simply need to set this as active configuration. In addition you will may to have to set up the bat file to suit the local GNU compiler – just open the bat file and modify the paths so that the GNU tools are found.

The bat file (Build\_LM3SXXXX.bat) can be found in `\Applications\uTaskerV1.4\GNU_LM3SXXXX` and can also be started by double clicking on it.

The GCC make file automatically generates a stand-alone as well as objects containing the µTasker 'Bare-Minimum' boot loader, plus an object suitable to transmission via web browser or FTP.

- `\Applications\uTaskerV1.4\GNU_LPC23XX\uTaskerV1.4.bin` – this is a stand-alone object file which can be loaded to the target board.
- `\Applications\uTaskerV1.4\GNU_LPC23XX \uTaskerBM.bin` – this is an object file with embedded boot loader which can be loaded to the target board.
- `\Applications\uTaskerV1.4\GNU_LPC23XX \z_Upload.bin` - this is an object file which can be posted to the target board using a web browser (when the boot loader version is installed)

*When compiling the project, expect neither warnings nor errors to be emitted by the GNU compiler.*

If your goal is a complete open-source tool chain then it is probably best to use the **Yagarto** solution. This is based on Eclipse as IDE and enables a complete solution from editing, building, debugging and version management. With low cost development boards and JTAG debugger from Olimex ([www.olimex.com](http://www.olimex.com)) and the famous tutorial from **James Lynch** about how to install, configure and run, professional projects are possible on a shoe-string. The stand-alone GNU build is compatible with this environment and more details about it can be found at <http://www.utasker.com/forum/index.php?topic=34.0> (including how to obtain the tutorial mentioned above).

#### Warning to users of limited compiler versions

*The full functionality of the uTasker demo project produces a code size of about 50k..60k, depending on the compiler used. This will be too large for 32k limited versions of commercial compilers and so it will be necessary to remove some functions, depending on which ones are interesting for you and which less. Later there is a comparison of code sizes required by the modules so that you can easily estimate what is possible and what needs to be trimmed out. Interesting projects are however still possible without exceeding the limits.*

## 5. Downloading the Code to the Target

Some of the development environments include support for downloading the code to the FLASH of the target so this procedure can be performed quite comfortably using a compatible JTAG debugger. If you own such a compiler and debugger you probably know well enough how this works.

Here I will look at a solution for programming code to FLASH using the in-built ISP boot loader in the NXP processor. This uses a freeware program called **FLASH Magic** which was developed by Embedded Systems Academy for NXP. The NXP devices contain a pre-programmed boot loader which enables software to be loaded via UART0. There are two possibilities to start the processor in this mode:

- either there is no valid code loaded (valid code is recognised by a reserved location in the interrupt vector table being equal to a check sum over the other vectors)
- or Port P2.10 is logic '0' at reset (note that the port varies between processor types but this is valid for the LPC2378 on the MCB2300/Olimex LPC2378-STK)

If you don't already have FLASH MAGIC you can get it from <http://www.flashmagictool.com/>

Once you have installed this you can load the code by following the following steps:

Step 1: Connect a COM port of your PC with a straight cable to COM0 of the MCB2300 or RS232\_0 on the Olimex LPC2378-STK

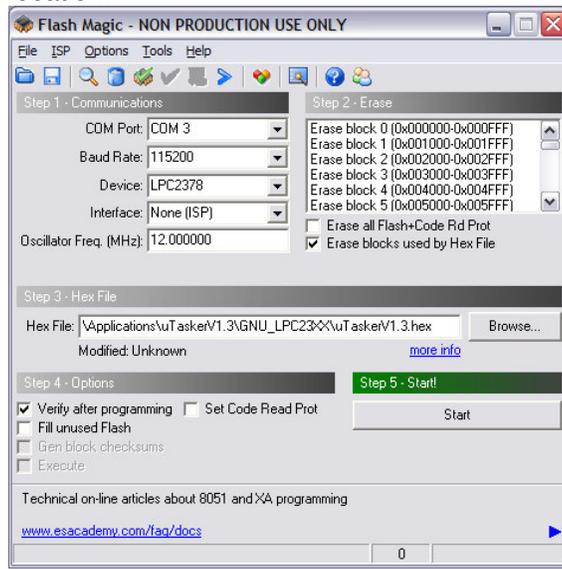
Step 2: Power on the evaluation board – its red power LED should light.

Step 3 – **MCB2300**: If your board already has code hold the INT0 button down while causing a reset with the RESET button. This will ensure that we are in ISP mode.

Step 3 – **Olimex**: Ensure that The ISP\_E and RST\_E jumpers are set on the board so that the FLASH Magic tool can control the reset and switching to the ISP mode via its RS232 control lines.

Step 4. Start FLASH Magic program and configure it as shown in the following screen shot. All settings are visible – set the COM to match yours and the hex file to match its

location.



Step 4. Click on the button “Start” to cause the program to connect to the boot loader, erase the FLASH and load the new code.

You should be done within a few minutes....

Are you surprised that the target behaves the same as the simulator? You shouldn't really be because that is exactly what the simulator is all about. It allows you to test your real code in real time and once it is working as you want it to, then you can transfer it to the real target. In fact you will find that your real target is not really necessary for most of your development work. Develop on the simulator and ship on the target – that is the way to do things really efficiently.

*Note that the I/O page allows the first 4 LEDs on the MCB2300 board to be controlled and the state of these to be displayed. Two LEDs are set as outputs by default. If P2.0 is reconfigured to be an input the blinking LED will no longer light.*

*The Olimex board has one LED which is connected to P1.21. This blinks by default when the software is running. If is configured as an input the blinking LED will be continuously lit*

## 6. Watchdog and Debugging

Before you get down to detailed work and debugging on the target there is one important topic which must be explained. It concerns the watchdog on the LPC23XX/LPC24XX.

The watchdog is a hardware timer which is used to monitor the correct operation of the processor. It has the capability that it can, when activated, detect a serious program error (unintentional forever loop or some other unrecoverable condition) and reset the processor so that it restarts and clears this otherwise serious fault condition.

The watchdog timer is configured in the µTasker demo project and is normally active and will cause a reset if it is not actively retriggered within a period of 2 seconds. Whether it actually runs and its exact characteristics can in fact be set up in the file `app_hw_lpc23xx.h`.

Here is the code which is responsible for its configuration. It can be found in `LPC23XX.c`:

```
#ifndef _DEBUG_CODE_
    INIT_WATCHDOG_DISABLE();
    if (!WATCHDOG_DISABLE()) {
        ACTIVATE_WATCHDOG();          // allow user configuration of internal watch dog timer
    }
#endif
```

The actual code which is inserted here is defined in the project specific set-up file `app_hw_lpc23xx.h`.

The following is an extract from the MCB2300 configuration:

```
#define INIT_WATCHDOG_DISABLE()    FIO2DIR &= ~(PORT2_BIT8);          // ensure input
#define WATCHDOG_DISABLE()        (! (FIO2PIN & PORT2_BIT8))      // pull this input down to
                                                                    disable watchdog
```

When the define `_DEBUG_CODE` is not active the watchdog will be enabled when the test is not true. The input P2.8 is configured as an input to enable this test – this is held to '0' to disable starting the watchdog.

The important point is that the watchdog will cause problems when debugging if it is enabled. It should thus be generally disabled when debugging, or else temporarily disabled via the optional disable input.

Furthermore it is important to understand the the watchdog itself is not reset when the watchdog timeout occurs – an active watchdog will remain active through a 'soft' reset. It is important that, when debugging, the watchdog is never enabled after a power on reset. If it is enabled once it is necessary to power cycle the board to successfully disable it again!

## 7. Specific Notes concerning the LPC2XXX

The μTasker project and the μTasker simulator are designed to work with all of the LPC2XXX devices – from the smallest LPC2101, with 8k FLASH and 2k SRAM up to the LPC24XX with integrated LCD controller and SD-RAM memory interface. The chip type can be easily selected in `config.h` and the project will automatically adapt itself where possible to include only features supported by the particular device.

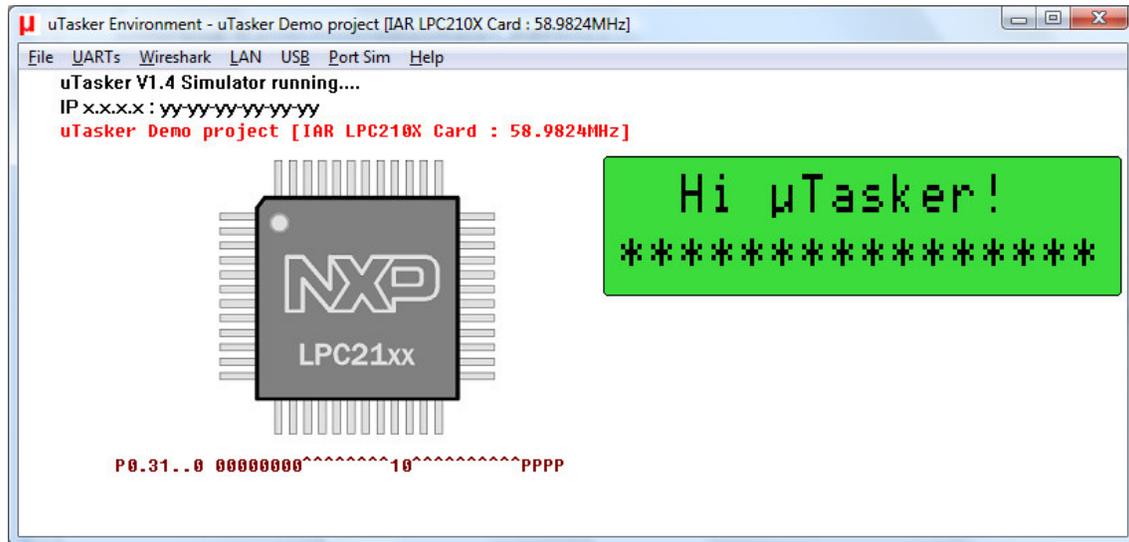


Figure 6. μTasker simulator showing the smaller LPC21XX in operation

The LPC23XX contains several RAM areas – one general purpose block, one designed in a way that it can be accessed optimally by the Ethernet controller and a similar one for the on-board USB controller. These RAM areas are not contiguous in RAM so only the general purpose one if used by the general code. The Ethernet block is used by the Ethernet driver and the USB block is presently not used.

The FLASH allows IAP (In Application Programming). This is used by the file system and the parameter system. The FLASH sector sizes are not all identical – some are 8k and some 32k and the 32k sectors are used mainly by the file system. Since this granularity is quite large, the μTasker file system uses a method of sub-files so that several files can be placed into one FLASH sector. For more details about this topic and the precautions required please see the documents in `\Applications\utaskerv1.4\WebPages\WebPagesLPC23xx`.

The large granularity principle (as used per default in LPC23XX projects) is described in <http://www.utasker.com/docs/STR91XF/FileSystemSTR91X.PDF> (this is for the STR91X but the principle is also valid for the LPC23XX).

## 8. A Tour through the Demo Web Server

Now that you have the possibility to simulate the demo and also run it on the target hardware it is time to take a more detailed look at what it does and then how it all works. The demo is designed to give you a practical platform to use as a basis for your own developments so it would be good to know how it can be best modified to suite your own needs.

### 8.1. Menu

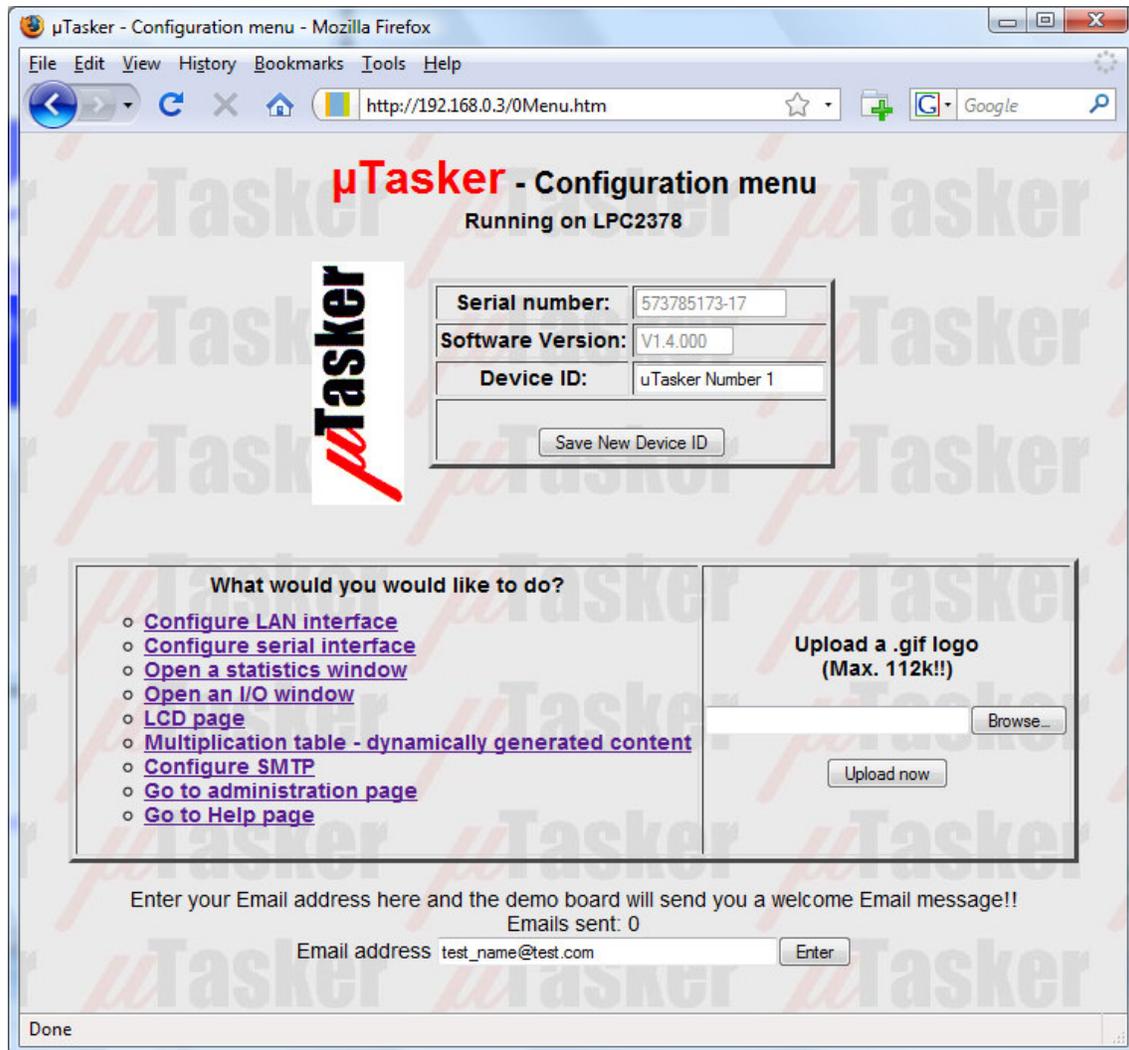


Figure 7. μTasker web server showing the menu page

**Serial number** – this is a decimal representation of the MAC address programmed in the device. Since the MAC address is normally unique it can serve also as a serial number if required. Notice that the serial number is displayed in grey and cannot be modified (disabled).

**Software version** – this is a string which is defined in `application.h`. It is also disabled so that it cannot be modified via the web side.

**Device ID** – this is a string defined in the `cParameters` structure in `application.c`. The default value can be set in code and also modified via the web page. When ‘*Save new device ID*’ is clicked, the new value is also saved to a parameter block in the internal flash so that it remains valid after the next reset or power up.

**Menu** – There are links to several other web pages allowing specific configurations to be displayed and modified.

**Email address** – If you have configured the demo project to support SMTP (`#define USE_SMTP` in `config.h`) then it is possible to enter an email address and send a test email to it. This entry is disabled if the support is not active in the project.

Note that it is necessary to configure some SMTP settings before an email can actually be sent – the default settings will not work since they are fictive and will not correspond with your email provider. Full details to the SMTP and its use in the µTasker demo project are given in the document <http://www.utasker.com/docs/uTasker/uTaskerSMTP.PDF>

**Logo and upload field** – The µTasker logo was uploaded to the file system as a GIF file when the files were loaded. It is possible to change this logo for a different GIF file by selecting a GIF and pressing the “*Upload now*” button. The GIF may be up to 112k in size.

## 8.2. LAN configuration

This side is a very critical side since it allows important Ethernet settings to be modified. The first time the device is started, the default MAC address – defined in the `cParameters` structure in `application.c` - is 0-0-0-0-0-0. It can be changed just once, after which it is displayed as disabled. This is because it is normally necessary to program a new device with a unique MAC address which should normally never be changed again. (*See the short discussion about MAC addresses at the bottom of this section for some details about how they are obtained.*)

The IP address, subnet mask and default gateway IP address can be modified as long as the device is not set to DHCP operation. If they are first modified and then the device is set to DHCP mode, they represent preferred settings to be requested during the DHCP procedure or for use as a last resort if no DHCP server is available.

The Ethernet speeds of 10M, 100M or auto-negotiation (full or half-duplex) can be set.

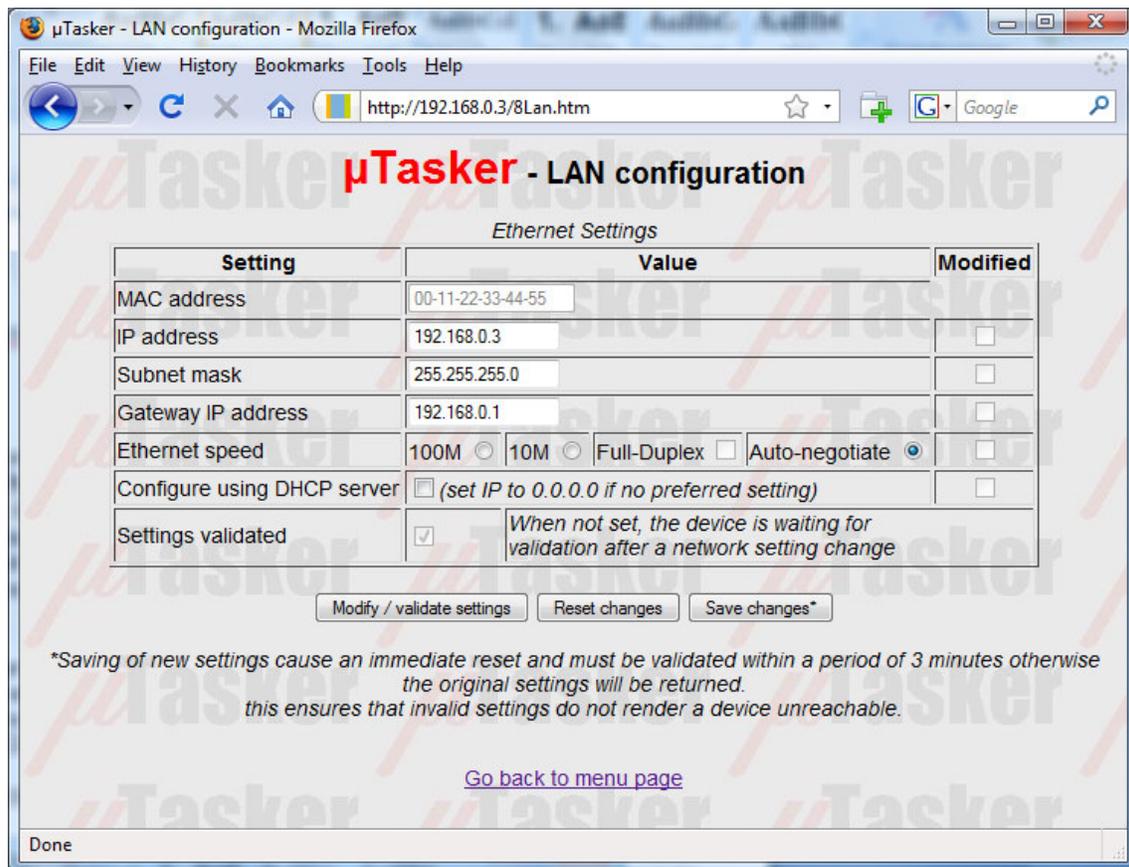


Figure 8. μTasker web server showing the LAN configuration page

Values, which are not disabled, can be modified and accepted by clicking on “Modify / validate settings” and changed values are indicated by a check box in the configuration table. They have not yet been committed to memory and the previous setting can be returned by clicking on “Reset changes”.

Once you are sure that the modifications are correct, they are saved to flash memory by clicking on “Save changes”. This will cause also a reset of the device after a delay of about 1

second and it will be necessary to establish a new connection to the device using its new settings and validate the changes.

***So why make it so complicated?*** Well it is simply to be sure that you have not modifying something which will render the device unreachable. For example, you may have changed its Ethernet speed from 100M to 10M although it is connected to a hub which only supports 100M and, to make matters worse, the device is not simply sitting on your office desk at arm's length but is at a customer's site several hours drive away. It would be a nasty situation if you had just lost contact with the equipment and have to somehow get it back on line although everyone at the customer's site has already left for a long weekend...

So this is where the validation part comes in. After the reset, the device sees that there are new settings in memory but also that these are only provisional (not validated) - the original values are also still available. These new values are nevertheless used for a 'trial-period' and a timer of three minutes duration started. It is now your job to establish a new web server connection to the device, using the new settings and to click on "Modify / validate settings". If you do this the check box names "Settings validated" will be checked, the new parameters are validated in flash and the old ones deleted; this in the knowledge that the new values are also really workable ones – how would you otherwise have been able to validate them?

Imagine however that a change really rendered the device unreachable – for example the LAN speed was really incompatible or you made a mistake with the subnet mask. After three minutes without validation, the provisional values are deleted from flash and, after a further automatic reset, the original settings are used again. After a short down period you can then connect as before the change and perform the modifications again, though correcting the previous mistake. There is therefore no danger of losing contact with the device, even when a 'silly' mistake is made.

### ***Some notes about MAC addresses***

If you are using your own device behind a router and it is not visible to the 'outside world' you can in fact program any MAC address that you like because it is in a private area. It just has to be unique in this private area. (This is also valid for a device sitting in a Demilitarized Zone- DMZ)

If however you are selling a product or the device is sitting directly on the Internet then it must have a world-wide unique MAC address which has to be purchased from IEEE. It is purchased either as a block (IAB) of 4k MAC addresses at a cost of about \$500, or if you are going to produce a large number of pieces of equipment you can purchase a unique company ID (OUI) of 16Million for about \$1'600 (plus \$2'000 if you don't want the OUI to be registered on the public listing).

It is then your responsibility to manage the assignment of these addresses in your own products.

The registration page is at: <http://standards.ieee.org/regauth/index.html>

For any one just wanting to make one or two pieces of equipment for hobby use it is a bit much to pay \$500 for a bunch of MAC addresses and use just one or two of them. Unfortunately it is no allowed to sell the rest on to people in similar situation because a block must always remain with the individual or organisation purchasing it.

One trick which is often used is to find out what the MAC address is in an old NIC from an old PC which is being scrapped. This MAC address is then used in your own piece of equipment and the old NIC destroyed. You can then be sure that the MAC address is unique and can not disturb when used for any imaginable application.

### 8.3. Serial Configuration

This web page allows the debug interface UART in the LPC2XXX to be set. See the UART User's guide for some more details of this interface:

<http://www.utasker.com/docs/uTasker/uTaskerUART.PDF>

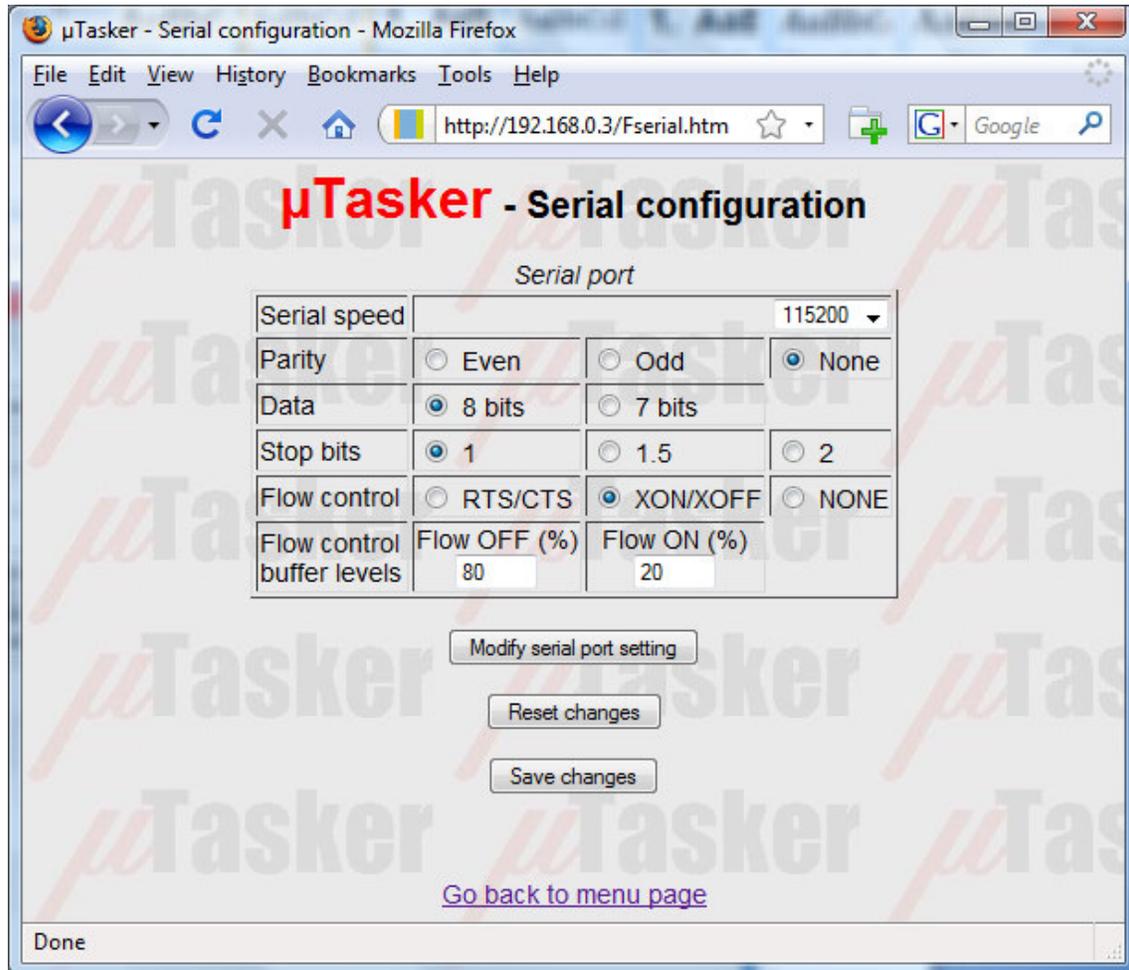


Figure 9. μTasker web server showing the serial interface configuration page

## 8.4. Statistics

The statistics window opens in a separate window and displays a table of the number of Ethernet frames counted by the Ethernet task. These are classified according to whether they were transmitted, received for the local IP address or received as broadcast frames and divided into IP protocol types. Since the values can continue to change, the user can update the window by clicking on the “Refresh” button. It is also possible to reset the counters if desired by clicking on “Reset frame stats”.

The contents of the ARP cache is also displayed and can be cleared by clicking on “Delete ARP entries”. These are equivalent to the DOS commands `arp -a` to display the local ARP cache and `arp -d` to delete it. Notice that when the ARP cache is deleted via the web page there will always remain one entry, this being the PC which commanded the deletion of the ARP table and updated the table. Due to the network activity it will always immediately be re-entered...

**μTasker - statistics**

Refresh

IP frame / ARP statistics

RX	Number of frames	TX	Number of frames	ARP entries in cache
<b>Total RX frames</b>	83	<b>Total TX frames</b>	67	192.168.0.100 00-1c-23-4d-59-e1
<b>Overruns</b>	0	<b>ARP frames</b>	1	-- --
<b>Frames to us</b>		<b>ICMP frames</b>	0	-- --
<b>ARP frames</b>	1	<b>UDP frames</b>	0	-- --
<b>ICMP frames</b>	0	<b>TCP frames</b>	66	-- --
<b>UDP frames</b>	0			
<b>TCP frames</b>	81			
<b>Broadcast frames</b>				
<b>ARP frames</b>	1	Reset frame stats		Delete ARP entries
<b>ICMP frames</b>	0			
<b>UDP frames</b>	0	<b>Other events</b>	0	

Done

Figure 10. μTasker web server showing the statistics page

## 8.5. I/O Page

The I/O window was discussed in step 12 of the tutorial and will not be discussed in any more detail here.

## 8.6. Administration Side

On this page the FTP server can be activated and deactivated. See also step 13 of the tutorial.

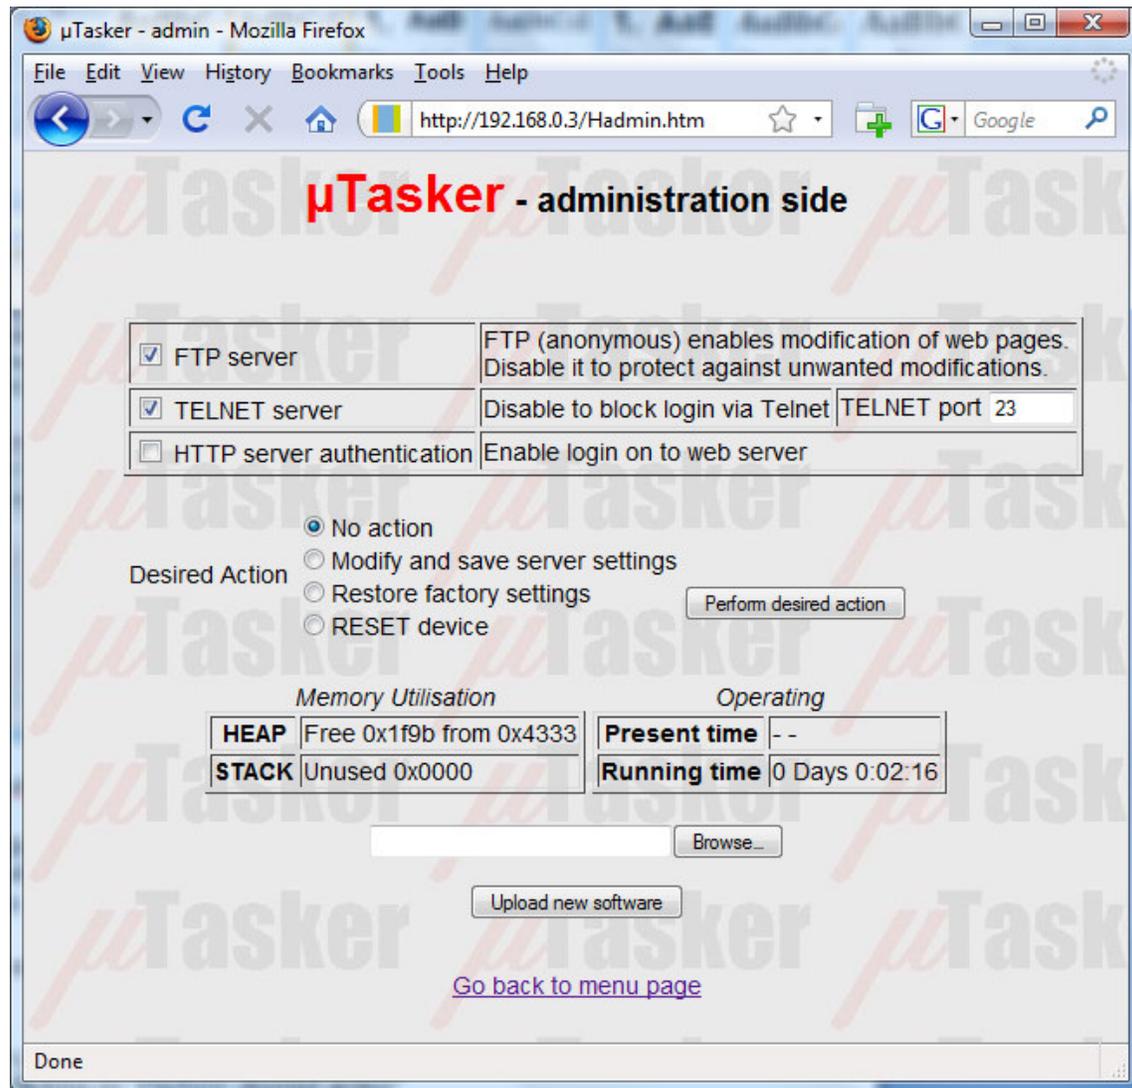


Figure 11. µTasker web server showing the administration page

The HTTP server can be configured to require basic authentication, meaning that it requests the user name and password when a new connection is established. These values are contained in the default settings in `application.c`.

A TELNET server can also be enabled or disabled, including setting it to a particular port number.

The previous three settings are modified and saved by setting the desired action to “*Modify and save server settings*” and clicking on “*Perform desired action*”.

A further action is “*Restore factory settings*” which returns the default settings as defined in `application.c`. This also provokes a reset of the device and if the default settings cause the network values to be modified it will automatically start a three minute validation period meaning that it will also be necessary to establish a connection using the ‘factory settings’ and validate them before the original ones will become permanent. This is again to ensure that there is no danger of losing contact with a remote device. Note that the MAC address is not reset since the Mac address should always remain unchanged in the device.

The last action is a simple reset of the device.

## Memory Utilisation

**HEAP** – the used heap and the maximum defined heap are displayed. Note here that the simulator will often display that more heap is used as needed in the real target due to the fact that some variables will be of different sizes when running on a PC compared to an embedded target. It allows the real use on the target to be monitored and the user can then optimise the amount of heap allocated to just support the worst case.

**STACK** – the amount of stack which has not been used (safety margin) is measured and displayed. As long as the value remains larger than 0 the system is not experiencing any memory difficulties. A value of zero indicates a critical situation and must be reviewed. The simulator however doesn’t perform this stack monitoring and will always display zero; its use is on the target where resources need to be monitored carefully.

The methods used for the memory monitoring are detailed in the µTasker main documentation.

**Present Time** – If the demo project is configured to collect the timer from a time server, this displays the present time. The time may also be available if the device which the µTasker is operating on processes a RTC.

**Running Time** - The time since the device started, or since the last reset, is displayed. This is useful when monitoring a remote device to ensure that it is indeed operating stably since a watchdog reset due to a software error would allow the device to remove but possibly not enable its occurrence to be readily detected.

## Upload new software

See the following µTasker document for further details about using the firmware upload capabilities of the µTasker project. This uses the HTTP post method to transfer new software to the target.

- [http://www.utasker.com/docs/uTasker/uTaskerBoot\\_003.PDF](http://www.utasker.com/docs/uTasker/uTaskerBoot_003.PDF)

## 8.7. LCD Side

When the graphical LCD support is enabled (define `SUPPORT_GLCD` in `config.h`). The LCD page will display the present LCD content. It also allows the user to post an image to be displayed on the graphical LCD.

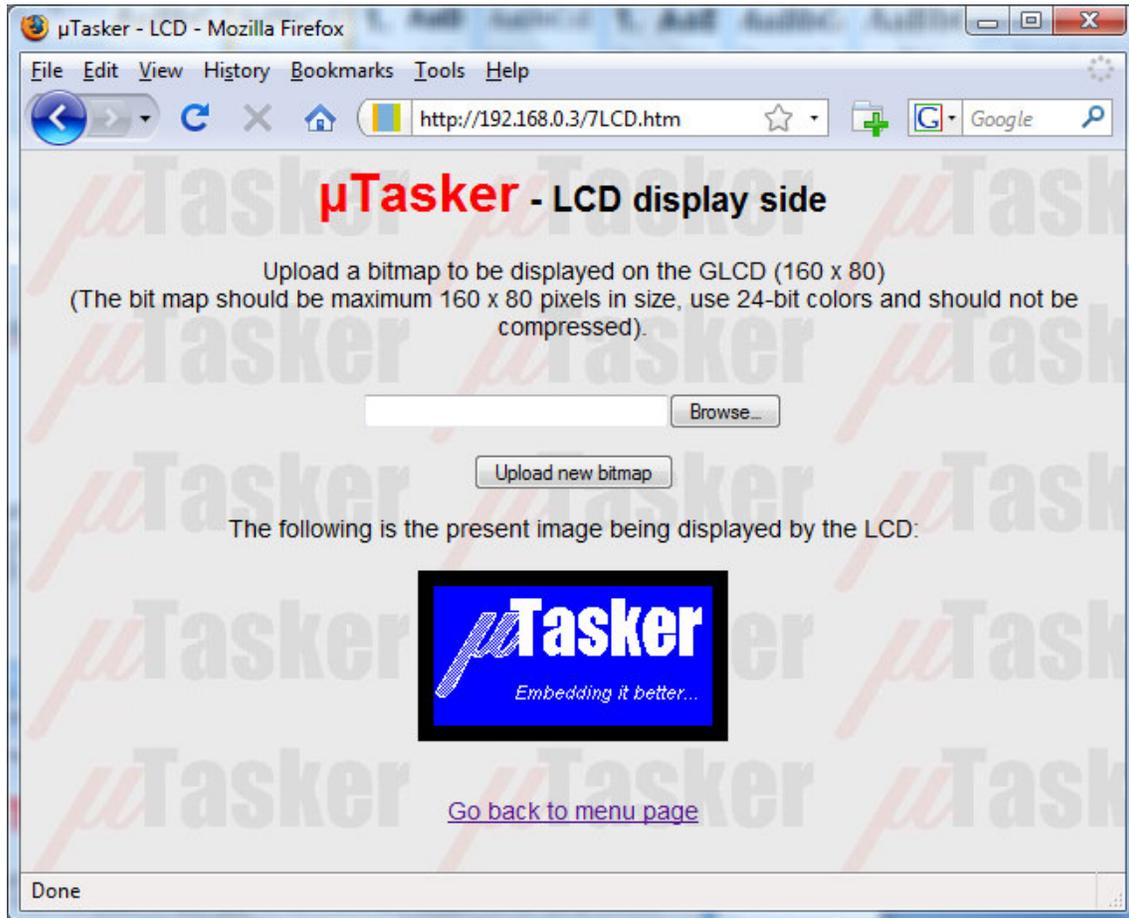


Figure 12. LCD page showing the present LCD image

For full details about the graphical LCD support in the project please consult the LCD User's Guide at <http://www.utasker.com/docs/uTasker/uTaskerLCD.PDF>

## 8.8. Email Configuration Side

The SMTP configuration can be performed on this web page corresponding to the email provider's requirements. For more details see the document

<http://www.utasker.com/docs/uTasker/uTaskerSMTP.PDF>

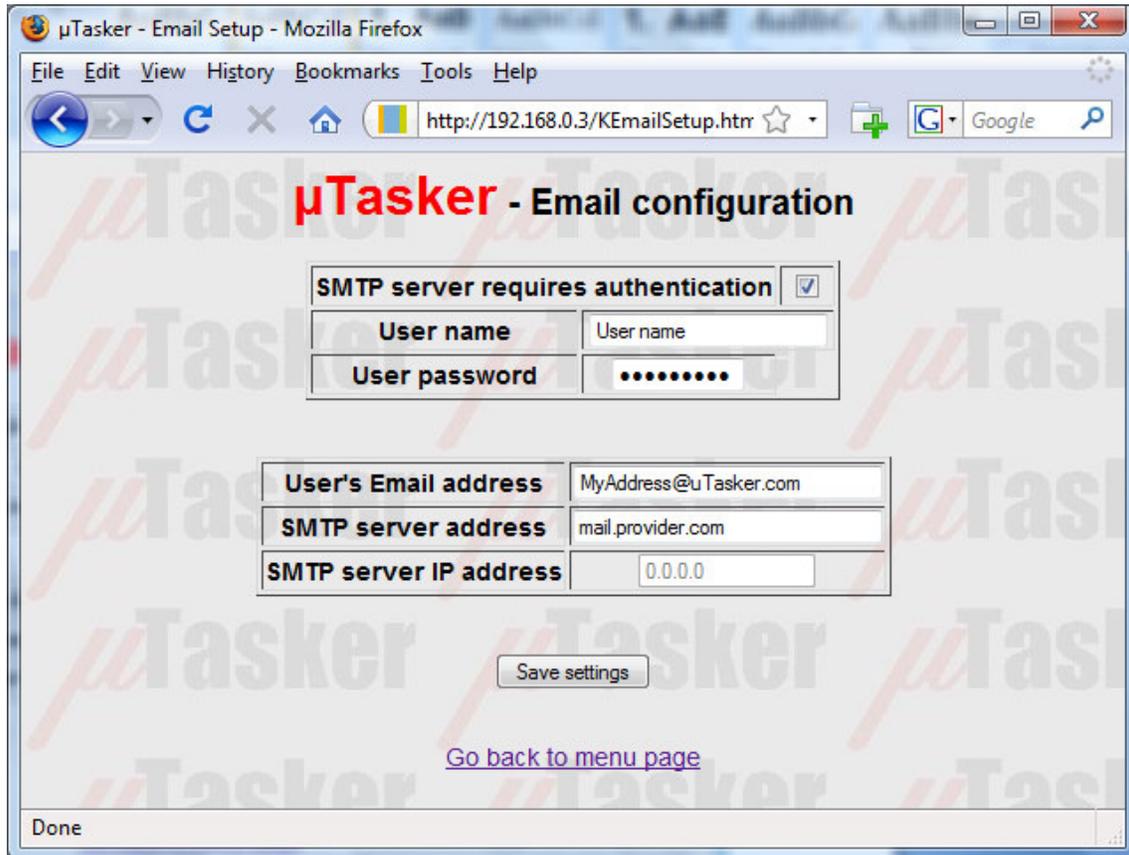
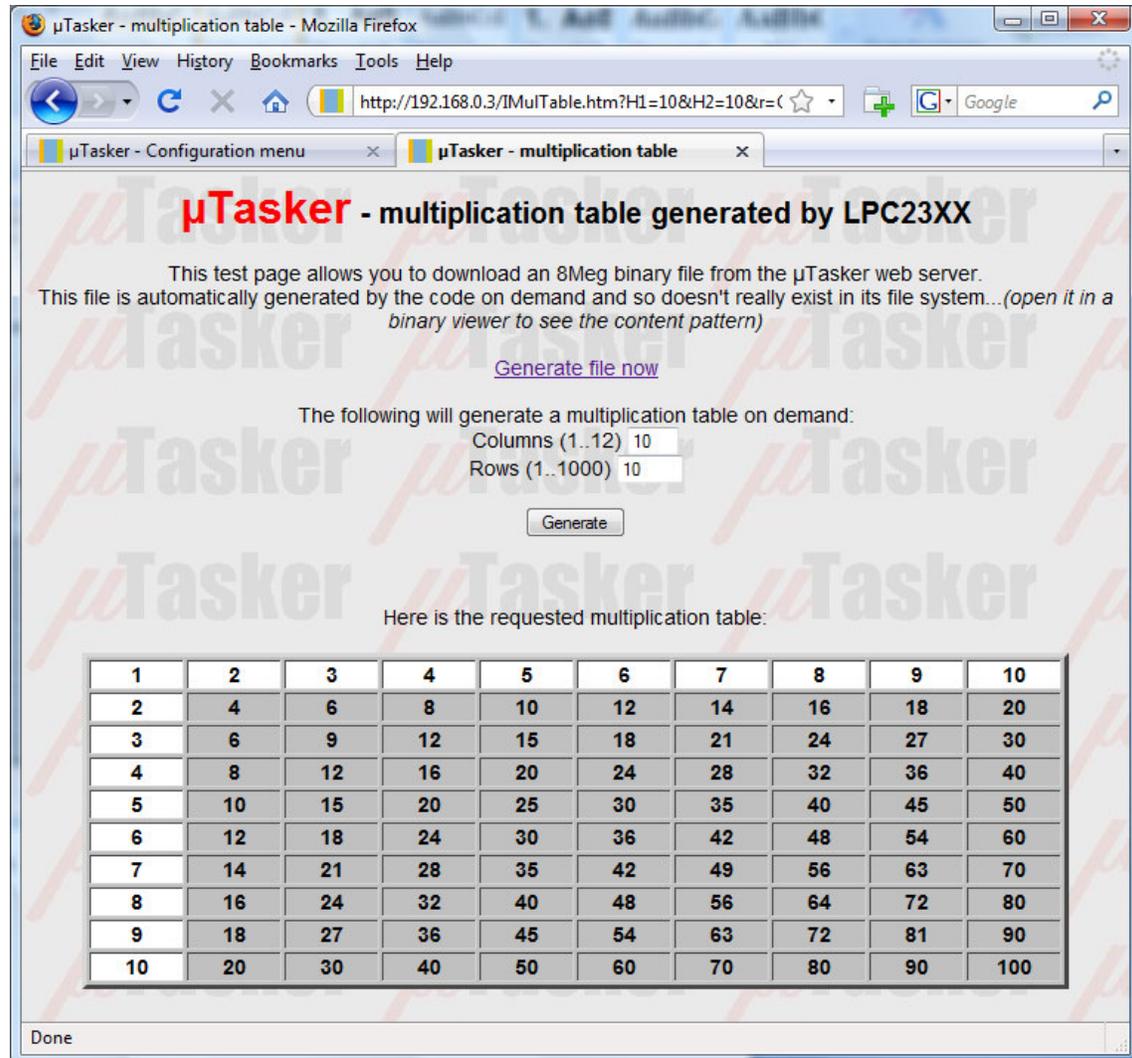


Figure 13. µTasker web server showing the Email configuration page

## 8.9. Dynamic Content Generation Side

The dynamic content generator web page demonstrates a code generated file and a multiplication table which can be generated on demand.

For full details of this operation see the document [http://www.utasker.com/docs/uTasker/uTaskerV1.3\\_Content\\_Generation.PDF](http://www.utasker.com/docs/uTasker/uTaskerV1.3_Content_Generation.PDF)



The screenshot shows a Mozilla Firefox browser window titled "µTasker - multiplication table". The address bar contains the URL "http://192.168.0.3/IMulTable.htm?H1=10&H2=10&r=C". The page content includes a title "µTasker - multiplication table generated by LPC23XX", a description of the test page, a "Generate file now" link, input fields for "Columns (1..12)" and "Rows (1..1000)" both set to "10", a "Generate" button, and a 10x10 multiplication table. The status bar at the bottom shows "Done".

**µTasker - multiplication table generated by LPC23XX**

This test page allows you to download an 8Meg binary file from the µTasker web server. This file is automatically generated by the code on demand and so doesn't really exist in its file system...*(open it in a binary viewer to see the content pattern)*

[Generate file now](#)

The following will generate a multiplication table on demand:

Columns (1..12)   
Rows (1..1000)

Here is the requested multiplication table:

1	2	3	4	5	6	7	8	9	10
2	4	6	8	10	12	14	16	18	20
3	6	9	12	15	18	21	24	27	30
4	8	12	16	20	24	28	32	36	40
5	10	15	20	25	30	35	40	45	50
6	12	18	24	30	36	42	48	54	60
7	14	21	28	35	42	49	56	63	70
8	16	24	32	40	48	56	64	72	80
9	18	27	36	45	54	63	72	81	90
10	20	30	40	50	60	70	80	90	100

Done

Figure 14. Dynamically generated multiplication table

## 9. Developing – Testing - Debugging

The demo project is designed to demonstrate a typical use of the operating system and TCP/IP stack. It is useful in itself as a starting point for many applications and this section looks at the support for developing, testing and debugging your own applications.

Until now you haven't actually had to develop anything since the project is delivered fully functional. To develop your own project it will be necessary to make configuration changes to suit your own use, to modify existing code and add new code of your own. To learn more about these aspects it is possible to study the documentation about the μTasker operating system and protocol stack. A more hands on approach is also possible by letting the demo project run and walking through the code parts which you are interested in – we did this briefly at the start of the demo but didn't linger to discuss any details. Here we will check out the advantages of the simulator by looking in more depth at a rather more complicated debugging session.

Debugging is performed for a number of reasons. It is a natural consequence of the test phase where unexpected program behaviour is experienced and the causes and reasons need to be understood before correcting the code. It is often also an integral part of the test phase itself when code reviews are performed, exception handling is to be exercised or software validation is required. The simulator allows a high degree of tests to be performed in comfort before going to the target testing phase, where such reviews would be rather more complicated.

So let's test something in the demo project. We'll test the simple PING ECHO utility which we already used once and we'll see how we can get to know the software in a very convenient and efficient manner. We will see how we can manipulate the operation to test and validate special cases and to make corrections in the code (and verify them too). First we will work ON LINE with the simulator, meaning that the simulator will be running effectively in real time and we will capture and analyse events. Afterwards we will see how to do the same OFF LINE, using a recording of the first case (which could also be a recording made when using a real target).

### ***Introducing Wireshark***

If you are using the μTasker then you will certainly be wanting to make use of its network capabilities and a tool to monitor network activity is essential. These are often called Network Sniffers since they can monitor, record and analyse network activity by watching what happens on the local Ethernet connection. The μTasker was designed with and around Wireshark (previously called Ethereal), a free and powerful network Sniffer. You can download this from <http://www.wireshark.org/>

If you haven't this program then don't delay – download it and install it and then we will get down to some work.

Now do the following steps:

1. Start the demo project simulation as described in the first tutorial.
2. Open a DOS window and prepare the PING command “ping 192.168.0.3”.
3. Start Wireshark and start monitoring the traffic on your local network (Capture | Start -> OK).
4. Enter return in the DOS windows so that the PING test is started.
5. Wait until the PING test has completed; there are normally 4 test messages sent.
6. Stop the Wireshark recording and save it to the directory  
Applications\uTaskerV1.4\Simulator\Ethereal\ with the name ping.eth

At this stage you can also look at the contents of the Wireshark recording and you will see something like the following:

No.	Time	Source	Destination	Prot	Info
1	0.437770	192.168.0.100	Broadcast	ARP	Who has 192.168.0.3? Tell192.168.0.100
2	0.439926	00:11:22_33:44:55	192.168.0.100	ARP	192.168.0.3 is at 00:11:22:33:44:5
3	0.439933	192.168.0.100	192.168.0.3	ICMP	Echo (ping) request
4	0.449898	192.168.0.3	192.168.0.100	ICMP	Echo (ping) reply
5	1.441345	192.168.0.100	192.168.0.3	ICMP	Echo (ping) request
6	1.441650	192.168.0.3	192.168.0.100	ICMP	Echo (ping) reply
7	2.442782	192.168.0.100	192.168.0.3	ICMP	Echo (ping) request
8	2.442915	192.168.0.3	192.168.0.100	ICMP	Echo (ping) reply
9	3.444225	192.168.0.100	192.168.0.3	ICMP	Echo (ping) request
10	3.444344	192.168.0.3	192.168.0.100	ICMP	Echo (ping) reply

It is also very possible that there are other frames on the network from other computers or your own computer talking with others, and it is possible to perform many filtering functions to remove these either from the visible display or from the recording file – just look in the Wireshark Help.

Here my PC has the local address 192.168.0.102 and initially doesn't know how to find the destination 192.168.03. An ARP resolve is sent and the simulator responds to it, informing at which MAC address its IP can be found at.

The ping test is repeated four times, each time receiving a reply from the simulator.

So let's see in some more detail how we can check the operation of this procedure and we will begin at the deepest point in the code, the receiving interrupt routine in the Ethernet driver, which is called when a complete Ethernet frame has been received.

In the VisualStudio project, locate the file `LPC23XX.c` (in `hardware\LPC23XX\`) and search for the Ethernet reception interrupt routine called `EMAC_Interrupt()` and put a break point in interrupt routine by positioning the cursor at the start of the routine and pressing F9.

It may be that you will find the code stopping at this break point before starting the ping test, which is probably because the code is receiving broadcast frames from your network. If this happens before you start the ping test just press F5 to let it run again (which you may have to repeat if there is a lot of activity...).

Repeat the ping test from the DOS window and the execution will stop at the breakpoint which you set in the interrupt routine. Since there will be no answer to the ping test as our code has been stopped, the ping test(s) will fail, but we don't worry about that since we have caught the event which we are going to use to analyse the flow through the driver, memory, the operating system and the ICMP routine. Now we can get to know the code in as much detail as we want.....and since we are working with the LPC23XX, this will also give us the opportunity to look at some of its internal registers on the way.

### A. Interrupt routine

```
// EMAC interrupt dispatcher
//
__interrupt void EMAC_Interrupt(void)
{
    while (EMAC_IntStatus & (EMAC_RxOverrunInt | EMAC_RxErrorInt | EMAC_RxFinishedInt |
EMAC_RxDoneInt)) { // read the interrupt status register
        if (EMAC_IntStatus & (EMAC_RxFinishedInt | EMAC_RxDoneInt)) {
            EMAC_IntClear = (EMAC_RxOverrunInt | EMAC_RxErrorInt | EMAC_RxFinishedInt |
EMAC_RxDoneInt);
            iInterruptLevel = 1; // ensure interrupts remain blocked
            fnWrite(INTERNAL_ROUTE, (unsigned char*)EMAC_RX_int_message, HEADER_LENGTH); //
inform the Ethernet task
            iInterruptLevel = 0;
        }
        else {
            EMAC_IntClear = (EMAC_RxOverrunInt | EMAC_RxErrorInt | EMAC_RxFinishedInt |
EMAC_RxDoneInt); // clear the interrupt status bit
        }
#ifdef _WINDOWS
        EMAC_IntStatus &= ~(EMAC_RxOverrunInt | EMAC_RxErrorInt | EMAC_RxFinishedInt |
EMAC_RxDoneInt); // clear interrupt flag when simulating
#endif
    }
    acknowledge_irq(); // require on exit of all irqs
}
```

The simulator has just received a frame for our MAC address or a broadcast address (the simulator doesn't disturb us with foreign MAC addresses when the EMAC in the LPC23XX has not be set up for promiscuous operation) and here we are in the interrupt routine. Now don't forget that this is the real code which will operate on your target and if this doesn't work properly it will also not work properly on your target.

We see that such routines are necessarily quite hardware specific. Here we see accesses to the internal register, `EMAC_IntStatus`. Search for this registers in the LPC23XX data sheet if you want to read exactly how it works; here is a very quick overview where we will also look at the registers in the simulated LPC23XX.

`EMAC_IntStatus` is the interrupt status register in the EMAC. We can search for it in our project by using "search in files", ensuring that the search path starts at the highest level in the µTasker project, and using the search files of type \*.c and \*.h.

It will be found in the code at several locations but also in the file `LPC23XX.h`, where it is defined as:

```
#define EMAC_IntStatus *(volatile unsigned long*) (ETHERNET_CONTROLLER_BLOCK +
0xfe0) // Interrupt Status Register (Read Only)
```

where `ETHERNET_CONTROLLER_BLOCK` is also defined locally twice:

```
#define ETHERNET_CONTROLLER_BLOCK ((unsigned char *) (&ucLPC23xx.ucEMAC))
#define ETHERNET_CONTROLLER_BLOCK 0xffe00000
```

On the target the register is located at long word address `0xffe00000` and when simulating it can be found in a structure called `ucLPC23XX`.

Double click on the word `ucLPC23XX` so that it becomes highlighted and then drag it into the watch window (Shift F9 opens a quick watch window). Expand the structure by clicking on the cross to the left of the name and you will see the various internal peripherals divided into sub-blocks. `EMAC_IntStatus` is in the sub-block `ucEMAC` so we also expand this sub-structure by clicking on the cross left of its name.

In the sub-block you will see all of the EMAC registers listed in the order in which they occur in memory as well as showing their present values. Step the code (F10) and you will see that an internal message is sent, defined by

```
static const unsigned char EMAC_RX_int_message[ HEADER_LENGTH ] = { 0, 0 ,  
TASK_ETHERNET, INTERRUPT_EVENT, EMAC_RX_INTERRUPT };  
// define fixed interrupt event
```

This internal message is sent to `TASK_ETHERNET` to inform that a frame has been received and is as such a method used by the operating system to wake up this Ethernet task. If you want to, you can also step into the write function to see the internal workings of how it wakes up the receiver task but I will assume here that all is working well and will just move to `TASK_ETHERNET` to see it actually receiving this interrupt message.

Before doing the next step, remove the break point in the interrupt routine by setting the cursor to its line and hitting F9.

## ***B. Receive Task***

Now open the file `ethernet.c` in the project directory TCP/IP and set the cursor to the bracket after the tasks routine `void fnTaskEthernet (TTASKTABLE *ptrTaskTable)`

With a right click, the menu item “Run to cursor” can be executed and you will see that this task is indeed started immediately since the program execution stops at this line.

Step slowly through the code using F10 and observe that the task reads the contents and interprets it as an `INTERRUPT_EVENT`. It calls `fnEthernetEvent()` to learn about the length of the received frame and to get a pointer to it – use F11 to step into this routine if you would like to see its internal workings, which I will not discuss here.

You will see that the pointer to the frame is called `rx_frame` and it is defined as a pointer to an `ETHERNET_FRAME`. Now, as we all know, a pointer is just a pointer – but the fact that it is defined as a pointer to a certain structure will help us greatly to interpret the contents of the received frame. We will first look at the frame as it is in memory and then how it looks in the debugger as a structure.

### C. Receive Frame

Let's start with simply looking at the raw data which we have just received from the Ethernet. Do this by placing `rx_frame` in a watch windows and expanding it so that the length of the frame (`frame_size`) and the pointer to the data (`ptEth`) are visible. `ptEth` has a value which you can double click on and insert to the clip board before pasting it into the address field of a memory display window (if this windows is not yet visible, activate it in the debug menu). Now you will see the raw Ethernet frame in memory – note that this is in the local computer's memory and the address of its location has nothing to do with the storage place on the real target. However its location is governed by the buffer descriptors in an analogue fashion to the operation in real device.

Now it is not exactly easy to understand the data but you should just be able to make out the MAC address at the beginning and the content of the ping test usually consists of `abcdefghijklmnop...` which is easily recognised. (Warning: it is possible that the frame which you actually captured is not the ping test but some other network activity, or even an ARP frame if the PC sending the ping had to re-resolve its MAC address. If this happens to be the case, set a break point at the location which the program is at and let it run again and hopefully it will come to this location the next time with the correct contents...).

Go back to the watch windows where the structure of `rx_frame` is being displayed and expand the sub-structure `ptEth`. This will start making life rather easier since it is showing us that the Ethernet frame is made up of a destination MAC address, a source MAC address, an Ethernet frame type and some data. Expand each sub-structure to see their exact contents, although it is not yet worth expanding the data field since we don't know what protocol its contents represent so it will not be displayed any better at the moment.

Name	Wert
rx_frame	{...}
frame_size	0x004a
ptEth	0x00456a40
ethernet_destination_MAC	0x00456a40 ""
[0x0]	0x00 ''
[0x1]	0x00 ''
[0x2]	0x00 ''
[0x3]	0x00 ''
[0x4]	0x00 ''
[0x5]	0x00 ''
ethernet_source_MAC	0x00456a46 ""
[0x0]	0x00 ''
[0x1]	0x0b ' '
[0x2]	0xdb 'Ù'
[0x3]	0xe6 'æ'
[0x4]	0x90 '¡'
[0x5]	0x2b '+'
ethernet_frame_type	0x00456a4c "I"
[0x0]	0x08 'I'
[0x1]	0x00 ''
ucData	0x00456a4e "E"

Adresse:	0x00457c1e
00457C1E	45 00 00 3C E...<
00457C22	9D 2D 00 00 ...
00457C26	80 01 1B DA ...Ù
00457C2A	C0 A8 00 66 A...f
00457C2E	C0 A8 00 03 A...'
00457C32	00 00 51 5C ...Q\
00457C36	02 00 02 00 ...
00457C3A	61 62 63 64 abcd
00457C3E	65 66 67 68 efgh
00457C42	69 6A 6B 6C ijkl
00457C46	6D 6E 6F 70 mnop
00457C4A	71 72 73 74 qrst
00457C4E	75 76 77 61 uvwa
00457C52	62 63 64 65 bcde
00457C56	66 67 68 69 fghi
00457C5A	41 43 41 43 ACAC
00457C5E	41 43 41 43 ACAC
00457C62	41 43 41 43 ACAC
00457C66	41 00 00 20 A...
00457C6A	00 01 00 00 ....
00457C6E	00 00 00 00 ....
00457C72	00 00 00 00 ....
00457C76	00 00 00 00 ....
00457C7A	00 00 00 00 ....
00457C7E	00 00 00 00 ....
00457C82	00 00 00 00 ....
00457C86	00 00 00 00 ....
00457C8A	00 00 00 00 ....
00457C8E	00 00 00 00 ....

## D. Frame protocol

Step slowly using F10 and you will see that the frame is checked for the ARP protocol, which it won't be if it is the ping request which we are analysing. It then calls `fnHandleIP()`, which we will enter using F11 since all such TCP/IP frames are built on the IP protocol.

After a couple of basic checks of IP frame validity, the pointer `received_ip_packet` is assigned to the data part of the received frame. Again this structure is very valuable to us since it can be dragged into the watch windows and now we suddenly understand how the IP fields are constructed. We can expand any field we want to see, such as the IP address of the source sending the IP frame. *Note that when looking at IP frames it may be worth requesting the debugger to display the contents in decimal rather than in hexadecimal by using right click and then selecting the display mode. The IP addresses are then better understandably and afterwards you can set the mode back to hexadecimal display since it is better for most other data fields.*

Name	Wert
received_ip_packet	0x00456a4e
version_header_length	69 'E'
differentiatedServicesField	0 ''
total_length	0x00456a50 ""
[0x0]	0 ''
[0x1]	60 '<'
identification	0x00456a52 "I0"
[0x0]	157 'I'
[0x1]	48 '0'
fragment_offset	0x00456a54 ""
[0x0]	0 ''
[0x1]	0 ''
time_to_live	128 'I'
ip_protocol	1 'I'
ip_checksum	0x00456a58 "IXA"
[0x0]	27 'I'
[0x1]	215 'x'
source_IP_address	0x00456a5a "A"
[0x0]	192 'A'
[0x1]	168 ''''
[0x2]	0 ''
[0x3]	102 'f'
destination_IP_address	0x00456a5e "A"
[0x0]	192 'A'
[0x1]	168 ''''
[0x2]	0 ''
[0x3]	3 'I'
ip_options	0x00456a62 "I"

Adresse: 0x00457c1e	Hex	ASCII
00457C1E	45 00 00 3C	E...<
00457C22	9D 2D 00 00	... ..
00457C26	80 01 1B DA	...U
00457C2A	C0 A8 00 66	A...f
00457C2E	C0 A8 00 03	A'...
00457C32	00 00 51 5C	..Q\
00457C36	02 00 02 00	... ..
00457C3A	61 62 63 64	abcd
00457C3E	65 66 67 68	efgh
00457C42	69 6A 6B 6C	ijkl
00457C46	6D 6E 6F 70	mnop
00457C4A	71 72 73 74	qrst
00457C4E	75 76 77 61	uvw
00457C52	62 63 64 65	bcde
00457C56	66 67 68 69	fghi
00457C5A	41 43 41 43	ACAC
00457C5E	41 43 41 43	ACAC
00457C62	41 43 41 43	ACAC
00457C66	41 00 00 20	A...
00457C6A	00 01 00 00	....
00457C6E	00 00 00 00	....
00457C72	00 00 00 00	....
00457C76	00 00 00 00	....
00457C7A	00 00 00 00	....
00457C7E	00 00 00 00	....
00457C82	00 00 00 00	....
00457C86	00 00 00 00	....
00457C8A	00 00 00 00	....
00457C8E	00 00 00 00	....
00457C92	00 00 00 00	....
00457C96	00 00 00 00	....
00457C9A	00 00 00 00	....
00457C9E	00 00 00 00	....
00457CA2	00 00 00 00	....
00457CA6	00 00 00 00	....
00457CAA	00 00 00 00	....
00457CAE	00 00 00 00	....

If you continue to step through the routine you will see that the IP frame is interpreted to see whether we are being addressed, it may cause our ARP cache to be updated and the checksum of the IP frame will also be verified. To return without stepping all the way you can use SHIFT F11 to return to the ETHERNET task code, where the protocol type is checked.

### ***E. ICMP Handling***

Note that each protocol type can be individually activated or deactivated in `config.h`. ICMP frames are only handled when the define `USE_ICMP` is set. If your project doesn't want to support ICMP then it can be simply removed...

Step into the ICMP routine (`fnHandleICP()`) by using F11. You will see that there is also a checksum in the ICMP field which is verified and the structure `ptrICP_frame` allows the ICMP fields to be comfortably viewed in the watch window.

Our frame should be of type `ECHO_PING`, which can also be individually deactivated in your project if you want to support ICMP but do not want the ping test to be replied to.

The received frame is sent back, after modifying a few fields, using the call `fnSendIP()`. I don't want to go into details about how the IP frame is constructed – you can see this in detail by stepping into the routine and observing what happens – but I should mention how the frame is sent out over the Ethernet since this is again a low level part which uses the LPC23XX registers again. Very briefly you should understand that the LPC23XX has been set up with a number of transmission buffers (`LAN_TX_BUF_COUNT` in `app_hw_LPC23XX.h`) into which the data is copied. The data is set up to respect the ICMP, IP and Ethernet layers as we observed in the received message and, once completely ready, the transmission is activated.

This transmission activation takes place in the file `LPC23XX.c` in the function `fnStartEthTx()` so search for it and set a break point there. Once the program reaches this point it will stop and you can see which registers are set up and verify that all is correct.

Here you will also see that the simulator comes into play once the data is ready and the registers have been set up correctly. This code, which will not be discussed here, basically tries to behave as the EMAC transmitter does by interpreting the register set up and transmitting the data buffer contents to your local network.

## 10. Off-Line Simulation

Normally the simulator, or the target, will have responded to the ping request very quickly and the ping test would have been successful. We, being human beings, are rather slow and we probably took several minutes to work our way through the code before your ping reply was finally sent back. This was of course much too late as the ping test has already terminated, informing that the test failed.

This is not only a problem with the ping test but with many protocols since they use timers and expect replies within quite a short time, otherwise an error is assumed and links break down. Debugging of such protocols can become quite hard work due to this fact.

This is where the μTasker can save the day again since it supports operation in OFF-LINE mode. This means simply that it can interpret Wireshark recordings as if they were real data from the network. Since we previously made a recording we can try this out right now!

### A. Prepare a break point

I suggest that you set a break point in the ICMP routine itself since we have already seen how the message arrives there and this will avoid false triggers due to broadcast frames in the mean time. Let the project run again by hitting F5.

### B. Open the Wireshark recording

In the menu `Wireshark` select `Load Wireshark/Ethereal file to play back`. Then select the `ping.eth` file which was previously recorded. If the simulation is repeated, the command `Replay last File` can be used instead – *this is only valid after the simulator is terminated normally*.

Perform this action and the Wireshark file will be interpreted. The times that frames arrived will be respected – if there is a gap or 1s between two frames then these two frames will also arrive with a 1s gap. You will see that the recorded ping will be received via the receive Ethernet interrupt routine and be passed up through the software until the breakpoint in the ICMP routine is encountered. The difference is that the break point and stepping in code can not disturb the protocol since also the recording is stopped.

This can be a great advantage when debugging protocols!

Before terminating this discussion there are a few points which should be noted, so here is a list of all relevant details which could be of use or interest.

1. When an Wireshark recording is played back, the internal NIC is closed so that there is no disturbance from the network.
2. The Wireshark recording can be repeated after it has terminated. There is no limit as to the number of times it can be repeated. Useful for incremental testing of a new piece of code...
3. After an Ethernet playback it is necessary to restart the simulator to use the simulator with the network again.
4. If you find a software error when stepping through the code, there is nothing to stop you making a correction without terminating the simulation session. Often VisualStudio can

recompile the new code and continue with the debugging session, thus allowing the correction to be immediately used – this is a major advantage of the VisualStudio environment! Try it and you will learn to love it...

5. Wireshark recordings do not have to originate from the simulator, they can be recordings from targets or from foreign devices. This means that recordings of a good known sequence can also be used as input to developments (for example the recording of an email being collected by your PC).

It is necessary that the simulator is set up to have the same IP and MAC addresses as the device in the recording and it will then receive all recorded frames to that device. Using this recording it is then possible to develop new code, verifying that it reacts as the original device did at each step.

Using this method, it is even possible to develop quite complicated protocols or services using a known good case as a reference. It can be basically tested before being switched onto the network so that there is a good chance that it will even work first go!

## 11. Conclusion

This tutorial first introduced the μTasker simulator's capabilities to execute the demo project software on a simulated LPC2XXX. It showed the basics of working with the simulator and its Ethernet capabilities before turning attention to running the same project on a target board.

By working with Wireshark and the μTasker simulator's ability to play back Wireshark recordings through its complete software code, interesting methods for increasing development, debugging and test efficiency have been demonstrated.

The LPC23XX project supports various other features. All features can be simulated using similar built-in techniques, thus enabling all aspects of real-world embedded projects to be performed with increased efficiency in comparison with convention development and debug techniques. Combining the μTasker simulator with the features of the μTasker demo project represents the fastest method of starting and completing embedded projects with the supported devices.

### *Change history:*

*0.01/19.01.2008: Provisional first draft.*

*0.02/24.01.2008: Add Olimex LPC2378-STK.*

*1.00/3.08.2009: Adapted for μTasker V1.4 with graphical LCD demonstration.*