

# XC-2 Ethernet Board Tutorial

---

*Version 10.4*

An interactive version of this tutorial is available in the XMOS Development Environment. To open, choose **Help ► Welcome** and select the XC-2 tutorial.



Publication Date: 2010/06/04

Copyright © 2010 XMOS Ltd. All Rights Reserved.

# 1 Introduction

The XC-2 is a development board based on the XMOS XS1-G4 device, intended for developing Ethernet-based products. It comprises a single XS1-G4 device, 10 LEDs, two press-buttons, an RJ-45 connector, JTAG and serial interfaces. A through-hole prototyping area is provided for connecting additional components.

The XS1-G4 consists of four XCores, each comprising an event-driven processor with tightly integrated general purpose I/O. Each processor provides up to eight threads, 400 MIPS and 64 KBytes of RAM. The I/O pins are connected to board components and can be directly controlled by software (see Section 10).

Programs are written using a combination of XC, C and C++. XC provides extensions to C that simplify the control over concurrency, I/O and time. These extensions map directly onto XCore hardware resources such as threads and ports, and are *efficient*—compiling into short instruction sequences, and *safe*—free from many sources of deadlock, race conditions and memory violations. This makes programs easy to write, understand and debug.

This tutorial provides an introduction to writing XC programs using the components integrated on your XC-2. It assumes that you are familiar with C [1]. This tutorial shows you how to:

- illuminate an LED on the board
- flash an LED at a fixed rate
- send the message "Hello World" to your PC over a serial link
- create multiple concurrent threads that flash LEDs at different rates
- send a token between multiple threads, each flashing an LED in sequence
- add a button listener thread that changes the token direction
- ping the IP address of your XC-2 from a PC, causing it to flash an LED

The example programs show you how XC simplifies the creation of programs for event-driven processors. Sections of the tutorial that introduce a new language feature include the new keywords or operators in their title.

The example programs have been tested with version 10.4 of the tools. Information on downloading, installing and using these tools is provided in a separate user guide [2].

## 2 Illuminate an LED: port, <:

This part of the tutorial shows you how to illuminate an LED on your XC-2, using an XC port and an output statement.

A port connects a processor to one or more physical pins and is used to interface with the board components. The port logic can drive its pins high or low, or it can sample the value on its pins.

The program below illuminates a single LED on your XC-2:

```
#include <platform.h>

out port x0ledB = PORT_LED_0_1;

int main() {
    x0ledB <: 1;
    while(1)
        ;
    return 0;
}
```

The declaration

```
out port x0ledB = PORT_LED_0_1;
```

declares an output port named `x0ledB`, which refers to the 1-bit port identifier `PORT_LED_0_1` (see Section 10). This identifier is defined in the board description file “XC-2.xn” and is exported to the header file `platform.h` during compilation.

Integrated input and output statements make it easy to express I/O operations on ports. In `main`, the statement

```
x0ledB <: 1;
```

outputs the value 1 to the port `x0ledB`, causing the port to drive its pin high. This in turn causes the LED connected to the pin to illuminate.

The port continues to drive the pin until instructed otherwise or until the program terminates. The infinite loop that follows the output statement prevents the latter condition.



Compile and run this program on your XC-2. The LED labelled X0LEDB should illuminate.

## 3 Flash an LED: timer, :>

This part of the tutorial shows you how to flash an LED at a fixed rate, using an XC timer and an input statement.

A timer is a hardware resource used to determine when an event happens or to delay execution until a particular time. Each timer contains a 32-bit counter that is incremented at 100MHz and whose value can be input at any time.

The program below flashes an LED on your XC-2:

```
#include <platform.h>
#define FLASH_PERIOD 2000000

out port x0ledB = PORT_LED_0_1;

int main(void) {
    timer tmr;
    unsigned ledOn = 1;
    unsigned t;
    tmr :> t;
    while (1) {
        x0ledB <: ledOn;
        t += FLASH_PERIOD;
        tmr when timerafter(t) :> void;
        ledOn = !ledOn;
    }
    return 0;
}
```

The statement

```
tmr :> t;
```

inputs the value of the timer `tmr`'s counter into the variable `t`. Having recorded the current time, the statement

```
t += FLASH_PERIOD;
```

increments this value by the required delay, and the statement

```
tmr when timerafter(t) :> void;
```

delays inputting a value until the specified time is reached. The processor must complete an input operation once a condition is met, even if the input value is not required. This is expressed in XC as an input to `void`.




Compile and run this program on your XC-2. The LED labelled `X0LEDB` should flash at the rate defined by the flash period.

## 4 Interface with a host over a serial link

This part of the tutorial shows you how to implement a UART protocol that transmits a message from the XS1-G4 to your PC over a serial link.

A UART translates data between parallel and serial forms for transmission over a serial link. Each bit of data is driven for a fixed period, during which time the receiver must sample the data. The diagram below shows the transmission of a single byte of data at a rate of 115200 bits/s:



The quiescent state of the link is high. A byte is sent by first driving a *start bit* (0), followed by the data bits and then a *stop bit* (1). A rate of 115200 bits/s means that each bit is driven for  $\frac{1}{115200} = 8.68\mu\text{s}$ .

The program below serializes a byte of data and transmits its individual bits over a 1-bit port using the UART transmission protocol:

```
#include <platform.h>

#define BIT_RATE 115200
#define BIT_TIME XS1_TIMER_HZ / BIT_RATE

void txByte(out port TXD, int byte) {
    unsigned time;
    timer t;

    /* input initial time */
    t :=> time;

    /* output start bit */
    TXD <: 0;
    time += BIT_TIME;
    t when timerafter(time) :=> void;

    /* output data bits */
    for (int i=0; i<8; i++) {
        TXD <: >> byte;
        time += BIT_TIME;
        t when timerafter(time) :=> void;
    }

    /* output stop bit */
    TXD <: 1;
    time += BIT_TIME;
    t when timerafter(time) :=> void;
}
```


The function `txByte` outputs a byte by first outputting a start bit, following by a conditional input on a timer that waits for the bit time to elapse; the data bits and stop bit are output in the same way.


The output statement in the `for` loop

```
TXD <: >> byte;
```

includes the modifier `>>`, which right-shifts the value of `byte` by the port width (1 bit) after outputting the least significant port-width bits. This operation is performed in the same instruction as the output, making it more efficient than shifting the value as a separate operation afterwards.

The XC-2 has a chip that performs a USB-to-serial conversion. When the board is connected to a PC using a USB cable, this chip presents a virtual COM port that can be interfaced using a terminal emulator.<sup>1</sup>

 Load a terminal emulator program on your PC and connect it to the virtual COM port provided by the XC-2. A simple terminal emulator is available from the XMOS community website.<sup>2</sup>

 Complete the program above by declaring a port for the UART and by writing a `main` function that outputs the message “Hello World!” to this port. You can find the relevant port in the diagram at the end of this tutorial on page 12. Compile and run this program on your XC-2. The terminal should receive and display the message.

## 5 Flash multiple LEDs concurrently: `par`, `on`

This part of the tutorial shows you how to flash multiple LEDs driven by different processors on your XC-2, using the XC `par` and `on` statements.

The `par` statement provides a simple way to create *concurrent threads* that can run independently of one another. The `on` statement instructs the compiler on which processor each port is connected and each thread is executed.

<sup>1</sup>Currently on MACs, the virtual COM port cannot be supported at the same time as the JTAG interface, preventing you from completing the following exercise.

<sup>2</sup>[http://www.xcore.com/tag\\_search?tag=uart](http://www.xcore.com/tag_search?tag=uart)

The program below creates four concurrent threads, each running an instance of a function that flashes an LED:

```
#include <platform.h>

#define FLASH_PERIOD 20000000

on stdcore[0] : out port x0ledB = PORT_LED_0_1;
on stdcore[0] : out port x0ledA = PORT_LED_0_0;
on stdcore[1] : out port x1ledB = PORT_LED_1_1;
on stdcore[1] : out port x1ledA = PORT_LED_1_0;

void flashLED(out port led, int delay);

int main(void) {
    par {
        on stdcore[0]: flashLED(x0ledB, FLASH_PERIOD);
        on stdcore[0]: flashLED(x0ledA, FLASH_PERIOD);
        on stdcore[1]: flashLED(x1ledB, FLASH_PERIOD);
        on stdcore[1]: flashLED(x1ledA, FLASH_PERIOD);
    }
    return 0;
}
```

The header file `platform.h` provides a declaration of the global variable `stdcore`, which can be used to specify the placement of port declarations and threads. The diagram at the end of this tutorial on page 12 shows which processor each LED is connected to.



An `on` statement may only be used with threads created by `main`, in which case `main` may contain only channel declarations, a single `par` statement and an optional `return` statement.



Modify the code from Section 3 to form the body of the function `flashLED`. Compile and run this program on your XC-2. Four LEDs on one side of the XS1-G4 should continually flash at a fixed rate.



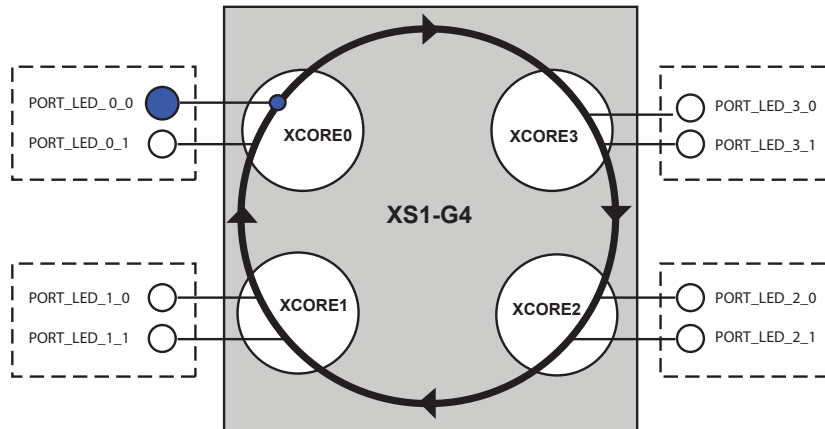
Experiment with different period values for each of the threads so that the threads can be seen to be operating independently of one another.



Extend this program to flash the LEDs on both sides of the XS-G4. Note that the LEDs on processor 3 are in reverse order on the XC-2.

## 6 Flash multiple LEDs in sequence: chan, chanend

This part of the tutorial shows you how to use XC channels to flash eight of the LEDs on your XC-2 in the round robin sequence illustrated below:



An XC channel provides a synchronous, bidirectional link between two threads. It consists of two channel ends, which two threads can use to interact on demand using the XC input and output statements.

The function below implements a component of the token ring illustrated above, repeatedly inputting a token from its left neighbor, flashing an LED and outputting the token to its right neighbor:

```
void tokenFlash(chanend left, chanend right,
               out port led, int delay, int isMaster) {
    timer tmr;
    unsigned t;
    if (isMaster) /* master inserts token into ring */
        right <: 1;
    while (1) {
        int token;
        left :> token; /* input token from left neighbor */
        led <: 1;
        tmr :> t;
        tmr when timerafter(t+delay) :> void;
        led <: 0;
        right <: token; /* output token to right neighbor */
    }
}
```

The first two function parameters are channel ends, the third an LED port and the fourth a Boolean value indicating whether or not the thread executing the function is the designated *master*. The master inserts a token into the ring.

The XC input and output statements are used to communicate the token between threads. As channels are synchronous, each output operation blocks until a matching input operation is ready, ensuring that precisely one thread has possession of the token at any time.

The function below constructs part of the token ring:

```
int main(void) {
    chan c0, c1, c2, ...;
    par {
        on stdcore[0]: tokenFlash(c0, c1, x0ledA, PERIOD, 1);
        on stdcore[0]: tokenFlash(c1, c2, x0ledB, PERIOD, 0);
        ...
    }
    return 0;
}
```

A channel is declared using the keyword `chan`. The locations of its two channel ends are established through its use in two statements of the `par`.

A total of eight channels are required to complete this program, each of which must be used in two threads: once as a left argument and once as a right argument to the function `tokenFlash`.



Complete this program, compile and run it on your XC-2. As the token cycles around the eight threads, the eight LEDs should each flash in sequence.

## 7 Add a button controller to the token ring: `select`

This part of the tutorial shows you how to detect a button press and respond to it, using the XC `select` statement.

A `select` statement is used to respond to one of a set of inputs, depending on which becomes ready first. If more than one of input becomes ready at the same time, only one is executed.

The function below waits for either a token to be received, in which case it passes it on, or for a button to be pressed, in which case it holds the token when received until the button is pressed again:

```

void buttonListener(chanend left, chanend right, port button) {
    int token;
    while (1)
        select {
            case left :> token :
                /* pass token on */
                right <: token;
                break;
            case button when pinsneq(0xf):> void :
                /* wait for token, then hold until button is pressed */
                left :> token;
                button when pinsneq(0xf):> void; /* push down */
                button when pinseq(0xf):> void; /* release */
                right <: token;
                break;
        }
}

```



The guarded input statement

```
case left :> token :
```

becomes ready when the token arrives, in which case it is input and passed to the next thread. The guard

```
case button when pinsneq(0xf):> void :
```



becomes ready when the value on the pins connected to the port `button` is not equal to the bit pattern `0xf`. This signifies that the button was pressed. The body of this case then waits for the button to be pressed again and released before continuing.

-  Add the function `buttonListener` to the token ring from the previous section that responds to either one of the two press-buttons. You can find their ports in the diagram at the end of this tutorial on page 12. Compile and run this program on your XC-2. Pressing one of the buttons should cause the LED to stop cycling after a full cycle is completed; pressing it again causes the next cycle to commence.
-  Modify the program so that pressing a button changes the cycle direction. The simplest solution is to add a `select` statement to the function `flashLED` that inputs from either its left or right neighbor, and extend `buttonListener` similarly.

## 8 Ping your XC-2 from a PC over its Ethernet connection

This part of the tutorial shows you how to integrate the XMOS Ethernet controller into your own design.

The 10/100-BASE-T Ethernet interface on the XC-2 is implemented using a standard physical layer interface device. The MAC and MII level protocols are implemented in software.<sup>3</sup>

-  Download the Ethernet software code, open the file “test.xc” and change the preprocessor value `OWN_IP_ADDRESS` to an IP address that your network can route. Compile and run this program on your XC-2. You should be able to ping and receive a response from the board. Note that you may need to disable your firewall.
-  Modify the function `demo` in the file “test.xc” so that it flashes one of the green LEDs each time it receives a ping request.

Congratulations, you’re now ready to start building Ethernet-based products using XC and your XC-2. For information on how to incorporate this software into your own design, refer to the files “README.txt” and “API.txt.”

## 9 What to read next

This tutorial provides only a basic introduction the XC language and XMOS architecture. The following documents provide more information on designing with XC on your XC-2 board:

- **Programming XC on XMOS Devices [3]**: Provides an in-depth tutorial on how to write XC programs, including case studies of real-world designs, along with the official XC language specification and details of its implementation on XS1 devices.
- **The XMOS Tools User Guide [2]**: Explains how to use the XMOS tools.
- **XC-2 Hardware Manual [4]**: Provides a functional description of the XC-2 board including the port-to-pin mappings of the through-hole prototyping areas, required to add additional components to your XC-2 based designs.

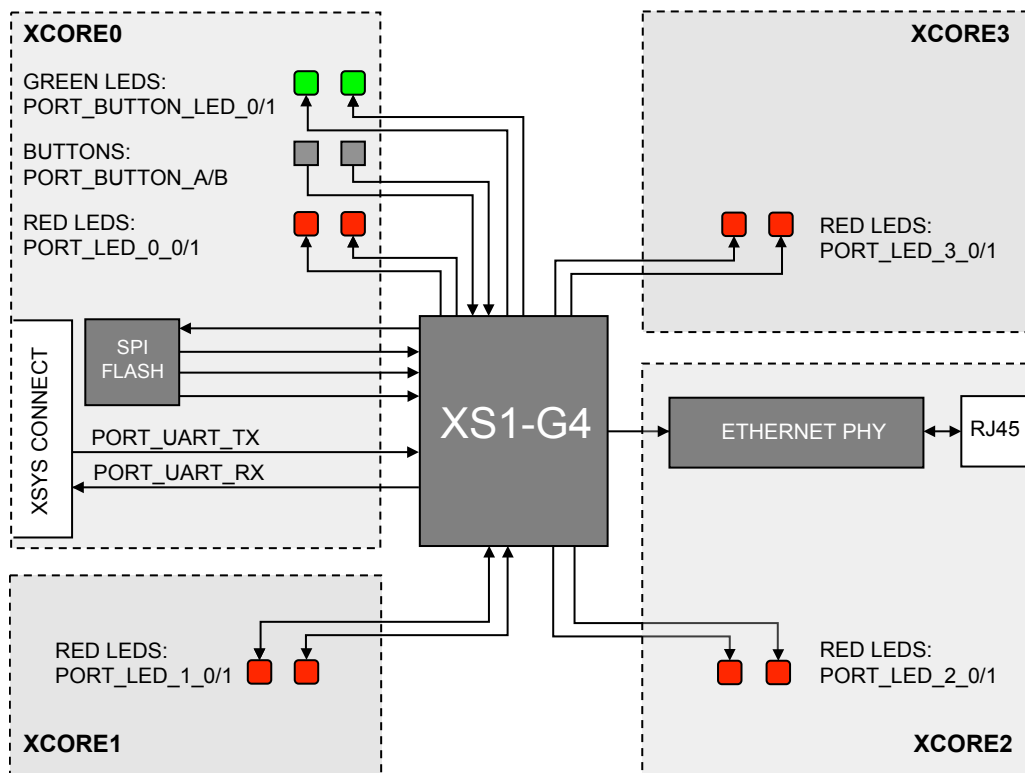
<sup>3</sup>The software can be downloaded as a ZIP file from <http://www.xmos.com/xc2/>. The ZIP file can be imported directly into Eclipse as a pre-configured project (for instructions see the Tools User Guide [2]) or uncompressed and built using the provided Makefile.

You may also find information from the following online resources useful:

- <http://www.xmos.com/>
- <http://www.xcore.com/>

## 10 XC-2 Port-To-Component Map

The diagram below shows how components on your XC-2 are connected to the processors on the XS1-G4, and gives the port names used in software.



## Bibliography

- [1] Brian W. Kernighan and Dennis M. Ritchie. *The C Programming Language*. Prentice Hall Press, Upper Saddle River, NJ, USA, 1988.
- [2] Douglas Watt and Huw Geddes. *The XMOS Tools User Guide*. XMOS Limited, 2009. [http://www.xmos.com/published/xtools\\_en](http://www.xmos.com/published/xtools_en).
- [3] Douglas Watt. *Programming XC on XMOS Devices*. XMOS Limited, Sep 2009. [http://www.xmos.com/published/xc\\_en](http://www.xmos.com/published/xc_en).
- [4] XMOS Limited. XC-2 Hardware Manual. Website, 2009. <http://www.xmos.com/published/xc2hw>.

## Disclaimer

XMOS Ltd. is the owner or licensee of this design, code, or Information (collectively, the "Information") and is providing it to you "AS IS" with no warranty of any kind, express or implied and shall have no liability in relation to its use. XMOS Ltd. makes no representation that the Information, or any particular implementation thereof, is or will be free from any claims of infringement and again, shall have no liability in relation to any such claims.

Copyright © 2010 XMOS Ltd. All Rights Reserved. XMOS and the XMOS logo are registered trademarks of XMOS Ltd in the United Kingdom and other countries, and may not be used without written permission. Company and product names mentioned in this document are the trademarks or registered trademarks of their respective owners. Where those designations appear in this document, and XMOS was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.