

## Advanced Computer Architecture

Lab2: High-Level Synthesis Lab

---

Johannes Kappes, Daniel Mueller-Gritschneider

18.11.2025

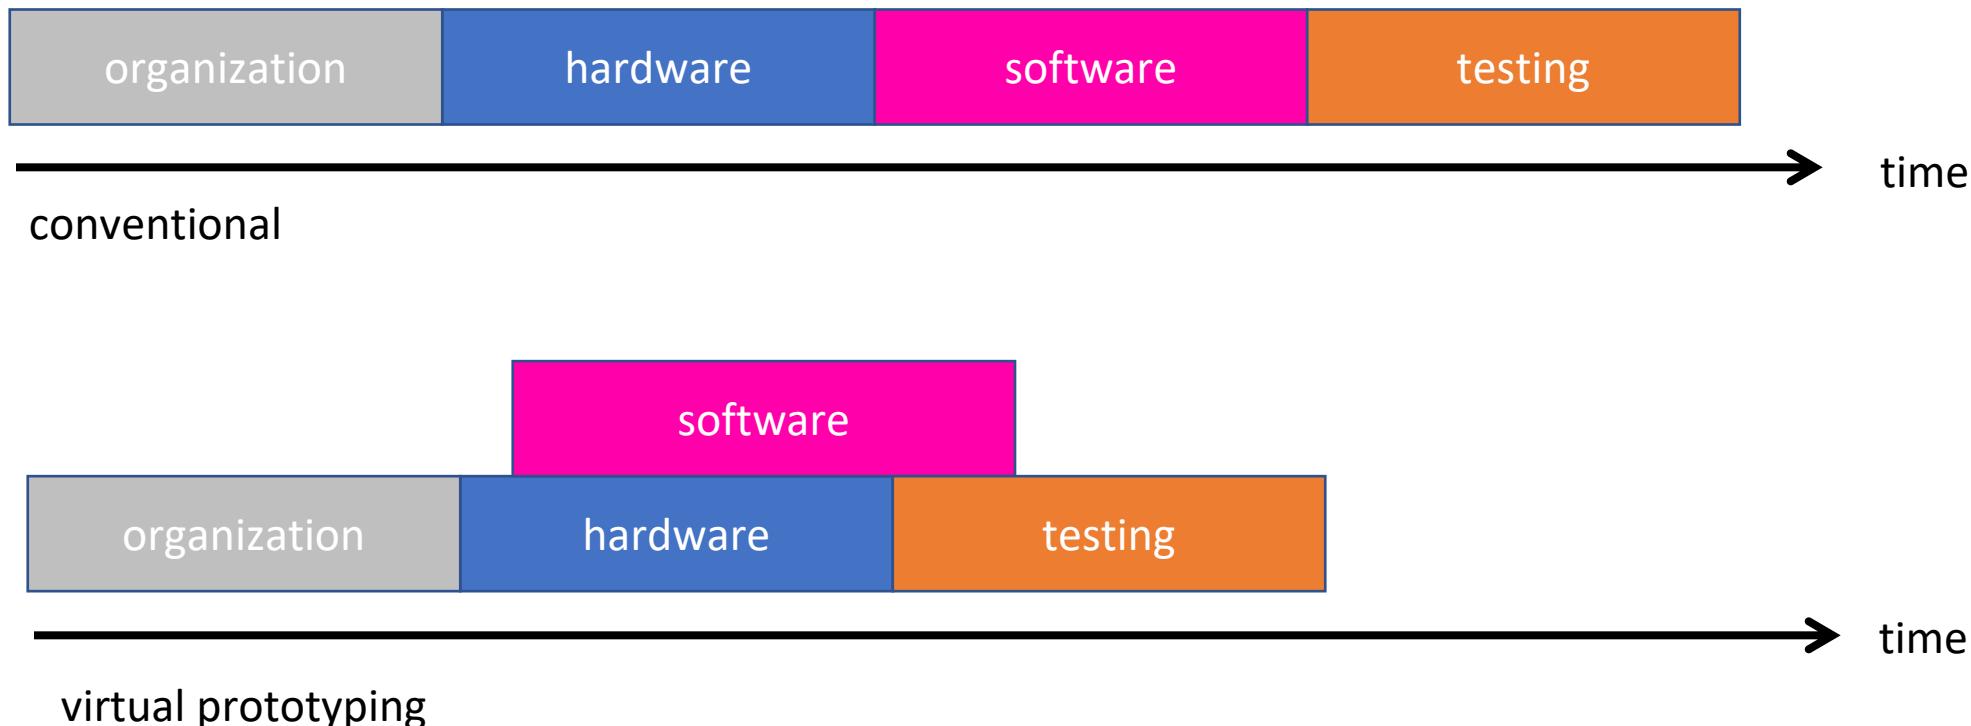
## 3 Modules:

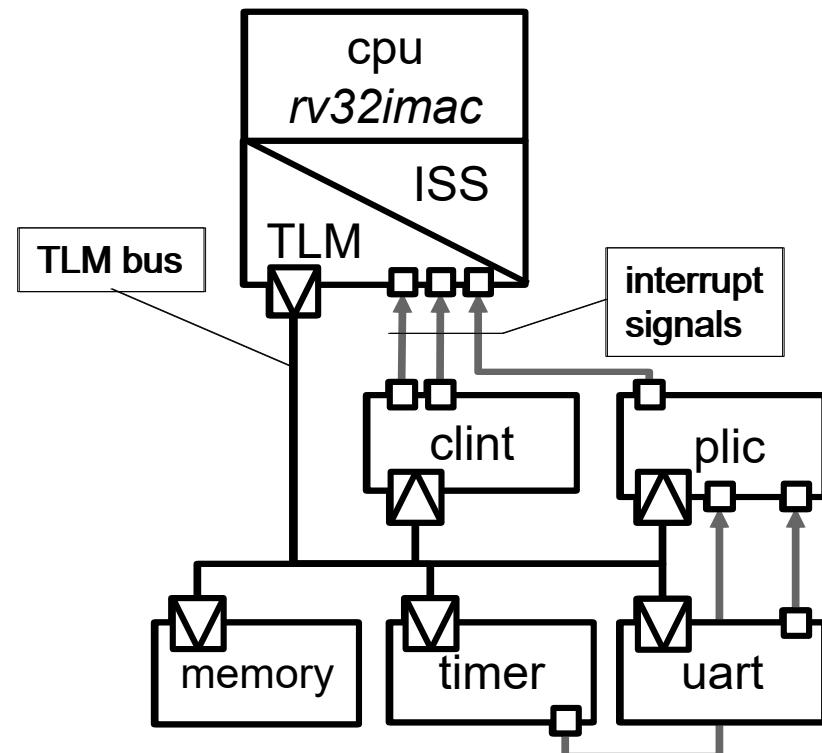
- M1 „Hello World!“ on a VP
- M2 Introduction into High Level Syntesis Tool (Vivado HLS)
- M3 Acceleration of Advanced Encryption Standard (AES) algorithm

# Module 1: Hello World! On a Virtual Prototype

# Module 1.1: Background

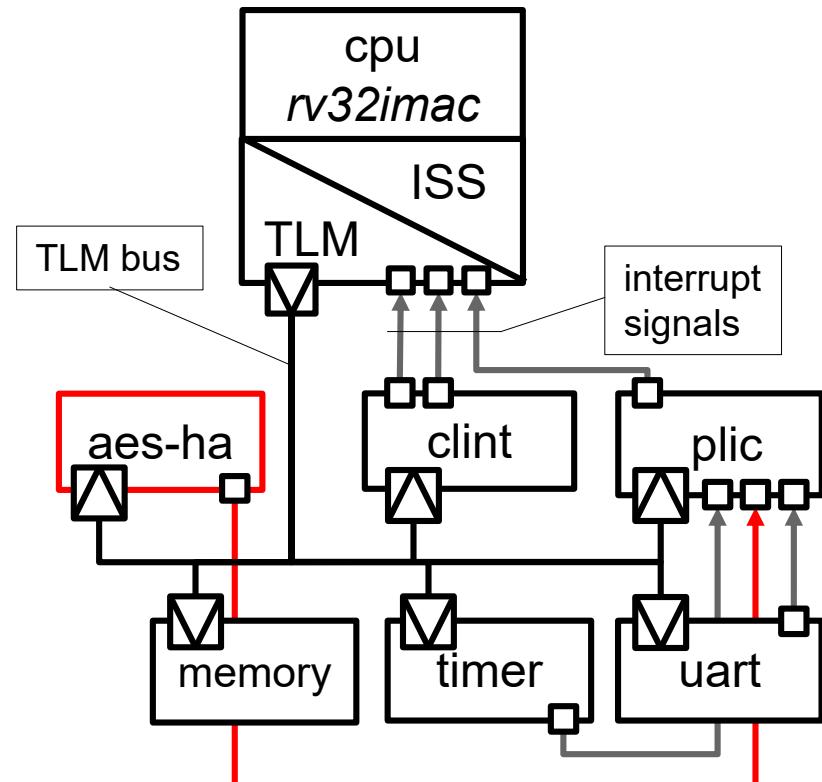
# Why Virtual Prototyping?





- 32-bit RISC-V CPU
- Core-local Interrupt Controller (clint) with Real-time Clock
  - (Machine Timer Interrupt)
  - MSI (Multi-Core)
- Platform-level Interrupt Controller (plic) as Interrupt MUX for:
  - UART with Interrupt
  - Timer with Interrupt
- Transaction Level Modeling (TLM) of Memory bus

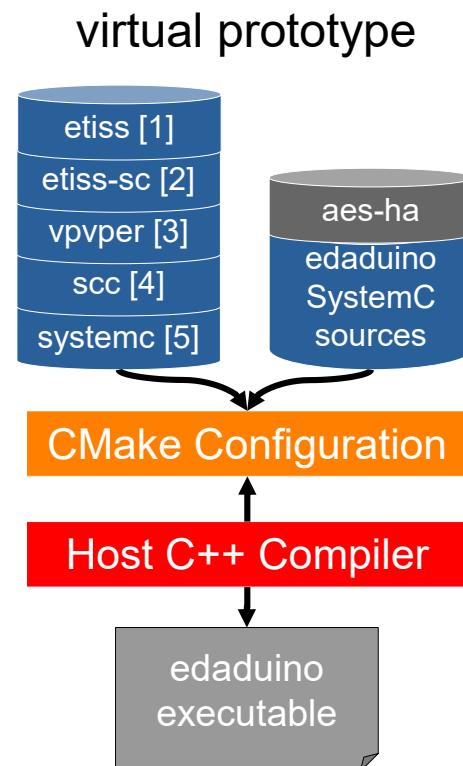
# EdaDuino - SoC (M3)



**- AES Hardware Accelerator  
(generated with Vivado-HLS)**

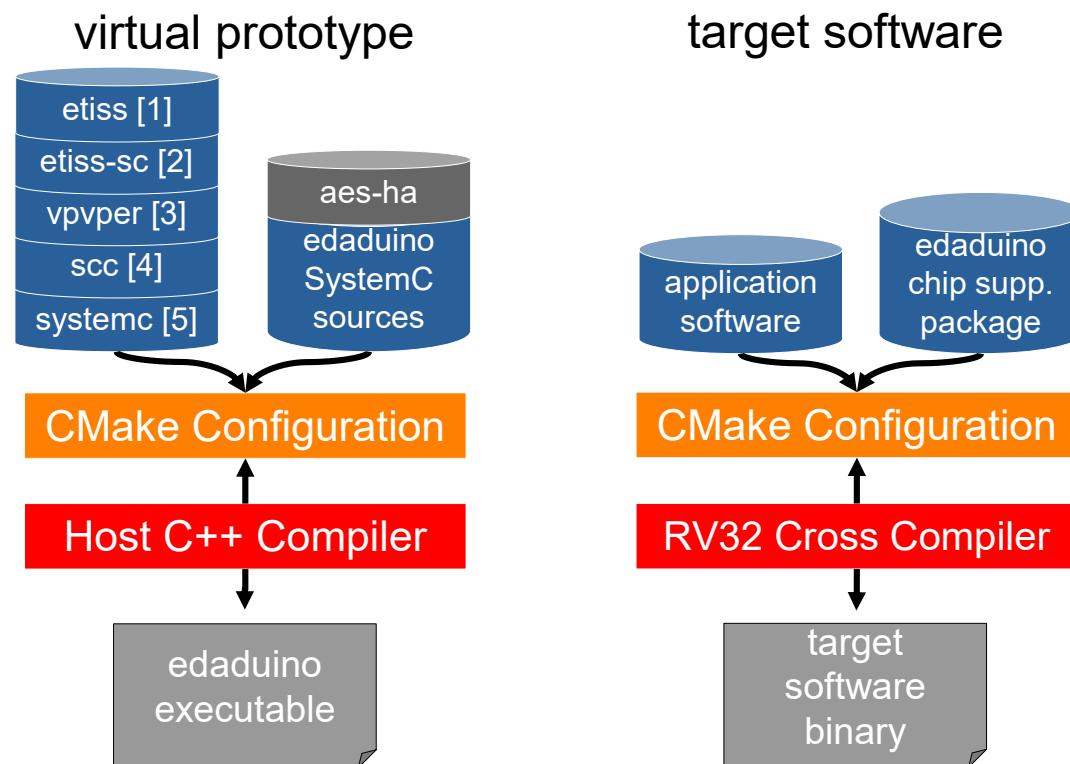
- 32-bit RISC-V CPU
- Core-local Interrupt Controller (clint) with Real-time Clock
  - (Machine Timer Interrupt)
  - MSI (Multi-Core)
- Platform-level Interrupt Controller (plic) as Interrupt MUX for:
  - UART with Interrupt
  - Timer with Interrupt
- Transaction Level Modeling (TLM) of Memory bus

# EdaDuino - Toolchain / Build System

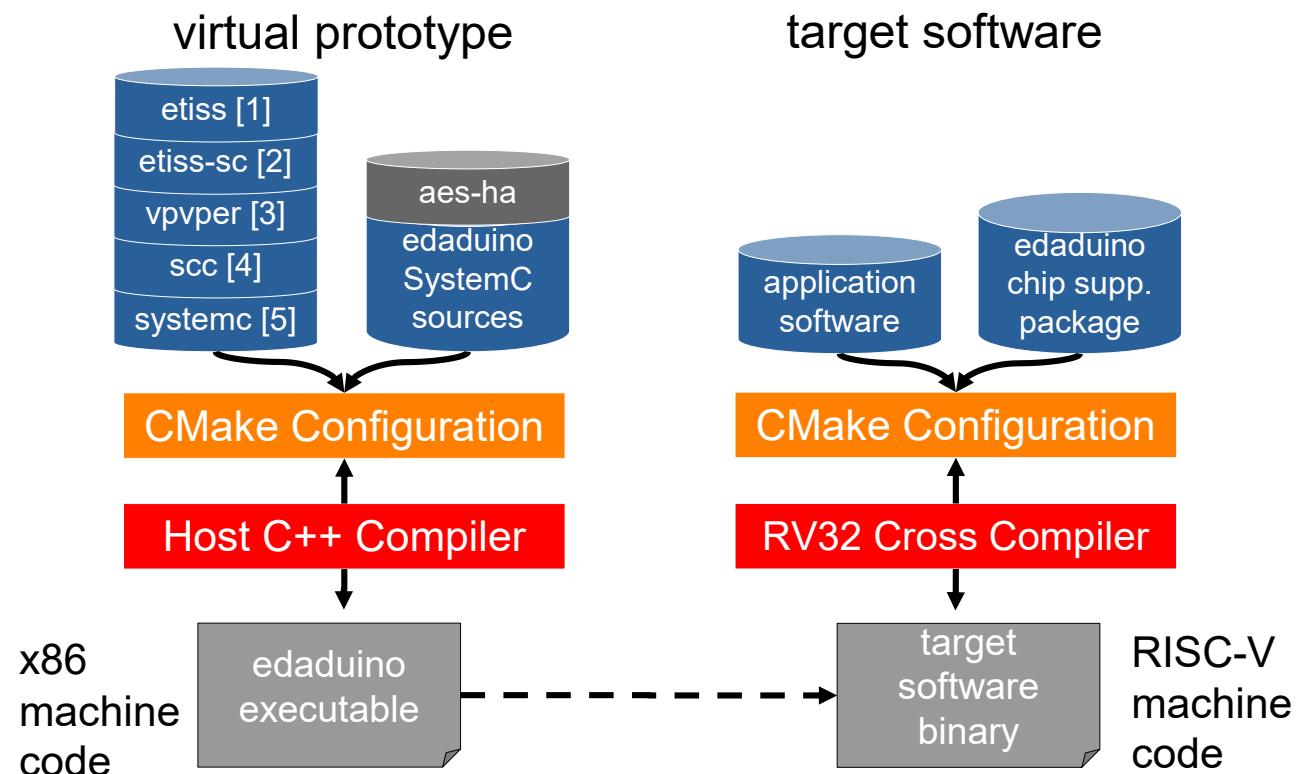


- [1] <https://github.com/tum-ei-eda/etiss>
- [2] <https://github.com/tum-ei-eda/etiss-sc>
- [3] <https://github.com/VP-Vibes/VPV-Peripherals>
- [4] <https://github.com/VP-Vibes/SystemC-Components>
- [5] <https://github.com/accellera-official/systemc>

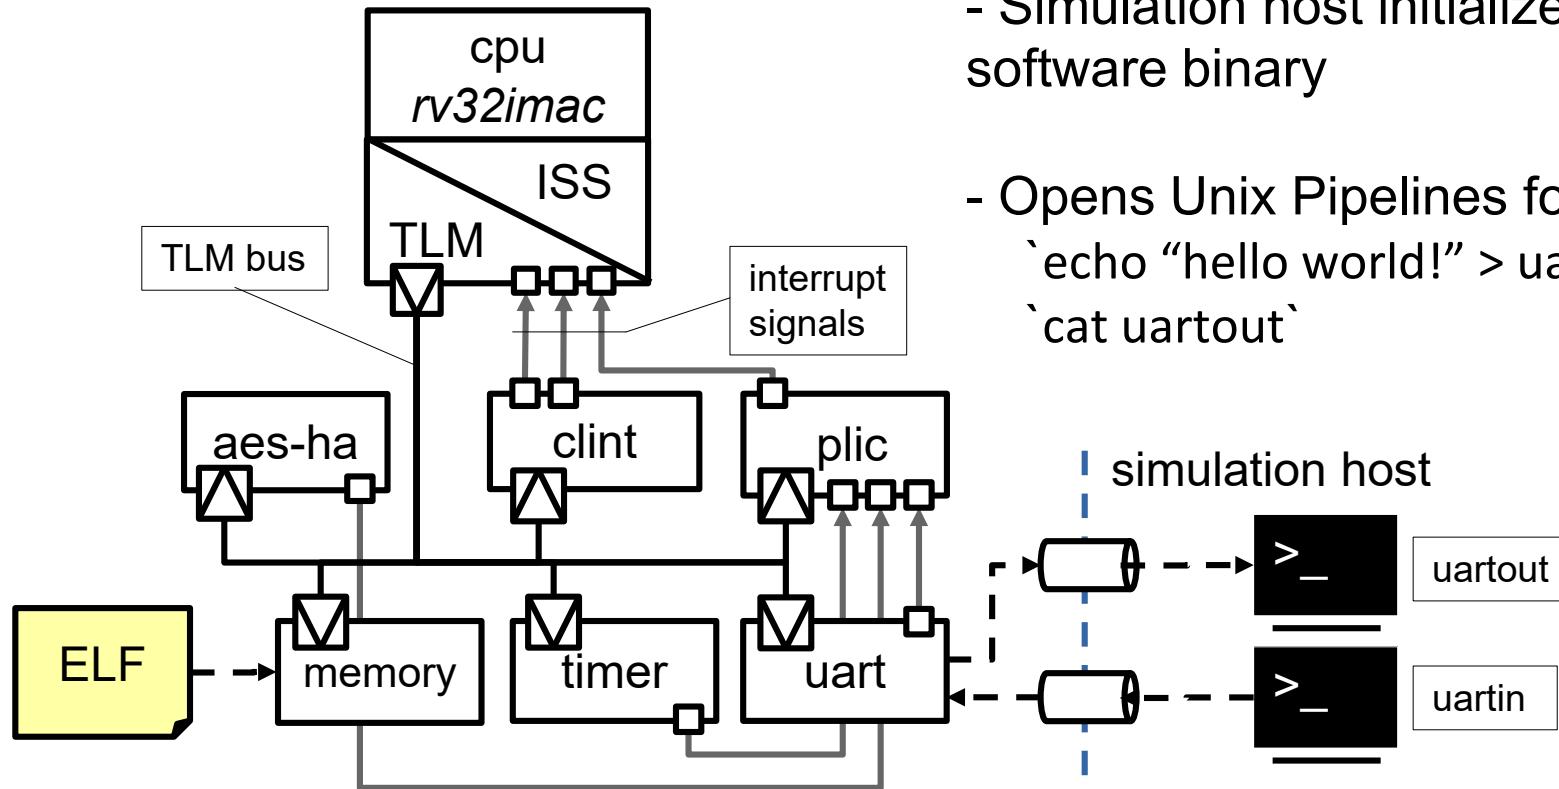
# EdaDuino - Toolchain / Build System



# EdaDuino - Toolchain / Build System



# EdaDuino - Virtual Platform



- Simulation host initializes memory with target software binary
- Opens Unix Pipelines for UART terminal, e.g.,  
`echo "hello world!" > uartin`  
`cat uartout`

## M1.2 Task: Utilizing UART Interrupt

# Polling vs. Interrupt-driven Hardware synchronization

- **Polling:**
  - cyclical sampling a peripherals status by an application program (synchronous)
  - cyclical checks can be scheduled
    - continuous sampling: “busy-wait”:
    - increase sampling period: better performance, risk of losing events
- **Interrupt-driven:**
  - peripheral reports a (configured) state change event to the CPU by an interrupt

# The System

- UART Peripheral \*
  - 16750 standard
  - memory mapped registers
  - single Interrupt line to CPU
- Software:  
“Serial Mirror”: UART received data is mirrored on Transmitter data

UART Register	Address std	Addr. abs EdaDuino	Type	Access Config
<b>RBR</b> <b>THR</b> <b>DLL</b>	base + (0*w)	0x1000_0000	<b>RO</b> <b>WO</b> <b>RW</b>	<b>LCR.DLAB = lo</b> <b>LCR.DLAB = lo</b> <b>LCR.DLAB = hi</b>
<b>IER</b> <b>DLM</b>	base + (1*w)	0x1000_0004	<b>RW</b> <b>RW</b>	<b>LCR.DLAB = lo</b> <b>LCR.DLAB = hi</b>
<b>IIR</b> <b>FCR</b>	base + (2*w)	0x1000_0008	<b>RO</b> <b>WO</b>	
<b>LCR</b>	base + (3*w)	0x1000_000c	<b>RW</b>	
<b>MCR</b>	base + (4*w)	0x1000_0010	<b>RW</b>	
<b>LSR</b>	base + (5*w)	0x1000_0014	<b>RO</b>	
<b>MSR</b>	base + (6*w)	0x1000_0018	<b>RO</b>	
<b>SCR</b>	base + (7*w)	0x1000_001C	<b>RW</b>	

\*elaborate guide to UART: <https://www.lammertbies.nl/comm/info/serial-uart>

# C Driver Library

sw/csp/inc/uart\_drv.h

```
/// \brief Uart control object. Holds states and registers
typedef struct Uart
{
    /* Registers */
    volatile uint8_t *rbr; ///< Receiver Buffer Register (make sure
    volatile uint8_t *thr; ///< Transmitter Holding Register (make
    volatile uint8_t *dll; ///< Divisor Latch (LSByte) Register (make
    volatile uint8_t *ier; ///< Interrupt Enable Register (make sure
    volatile uint8_t *dlm; ///< Divisor Latch (MSByte) Register (make
    volatile uint8_t *iir; ///< Interrupt Identification Register
    volatile uint8_t *fcr; ///< FIFO Control Register [WO]
    volatile uint8_t *lcr; ///< Line Control Register [RW]
    volatile uint8_t *mcr; ///< Modem Control Register [RW]
    volatile uint8_t *lsr; ///< Line Status Register [RO]
    volatile uint8_t *msr; ///< Modem Status Register [RO]
    volatile uint8_t *scr; ///< Scratch [RW]
    /* Callbacks */
    void (*irq_callback)(void); //< Function pointer to custom int
} uart_t;

void uart_set_dlab(uart_t *handle);
void uart_reset_dlab(uart_t *handle);
void uart_set_IERbit(uart_t *handle, uint8_t nbit);
void uart_reset_IERbit(uart_t *handle, uint8_t nbit);
void uart_set_irq_callback(uart_t *handle, void (*fptr)(void));
void uart_enable_interrupt(uart_t *handle, ier_bits_t interrupt);
void uart_disable_interrupt(uart_t *handle, ier_bits_t interrupt);

void uart_send_char(uart_t *handle, const char c);
char uart_get_char(uart_t *handle);
```

UART Register	Address std	Addr. abs. EdaDuino	Type	Access Config
<b>RBR</b> <b>THR</b> <b>DLL</b>	base + (0*w)	0x1000_0000	<b>RO</b> <b>WO</b> <b>RW</b>	<b>LCR.DLAB = lo</b> <b>LCR.DLAB = lo</b> <b>LCR.DLAB = hi</b>
<b>IER</b> <b>DLM</b>	base + (1*w)	0x1000_0004	<b>RW</b> <b>RW</b>	<b>LCR.DLAB = lo</b> <b>LCR.DLAB = hi</b>
<b>IIR</b> <b>FCR</b>	base + (2*w)	0x1000_0008	<b>RO</b> <b>WO</b>	
<b>LCR</b>	base + (3*w)	0x1000_000c	<b>RW</b>	
<b>MCR</b>	base + (4*w)	0x1000_0010	<b>RW</b>	
<b>LSR</b>	base + (5*w)	0x1000_0014	<b>RO</b>	
<b>MSR</b>	base + (6*w)	0x1000_0018	<b>RO</b>	
<b>SCR</b>	base + (7*w)	0x1000_001C	<b>RW</b>	

# The Task

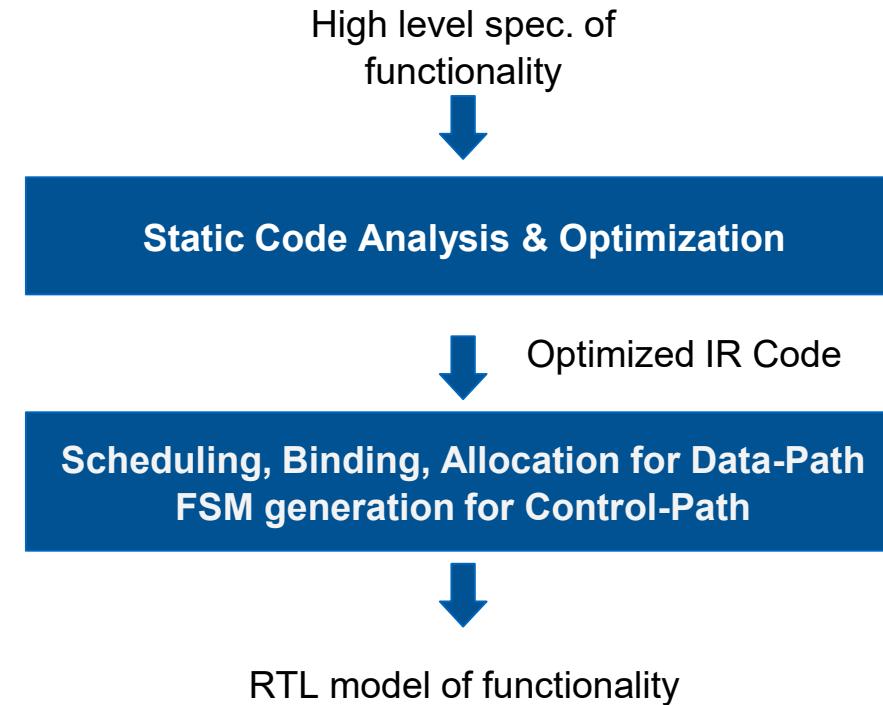
1. Write a `uart\_send\_string(const char\* str)` driver function
2. Check the pre-implemented busy-wait “Serial Mirror”
3. Replace the busy-wait poll with a wait for interrupt (WFI)
4. Write an **UART\_IRQ\_HANDLER**
5. Check your interrupt-driven solution, if it still performs the “Serial Mirror”

# Learning Goals?

- Basics of (embedded) Software Compilation Flow
- Programming Low-Level Memory Mapped I/O
- Basics of Bare-metal programming
- Basics of ETIIS VP for Module part 3

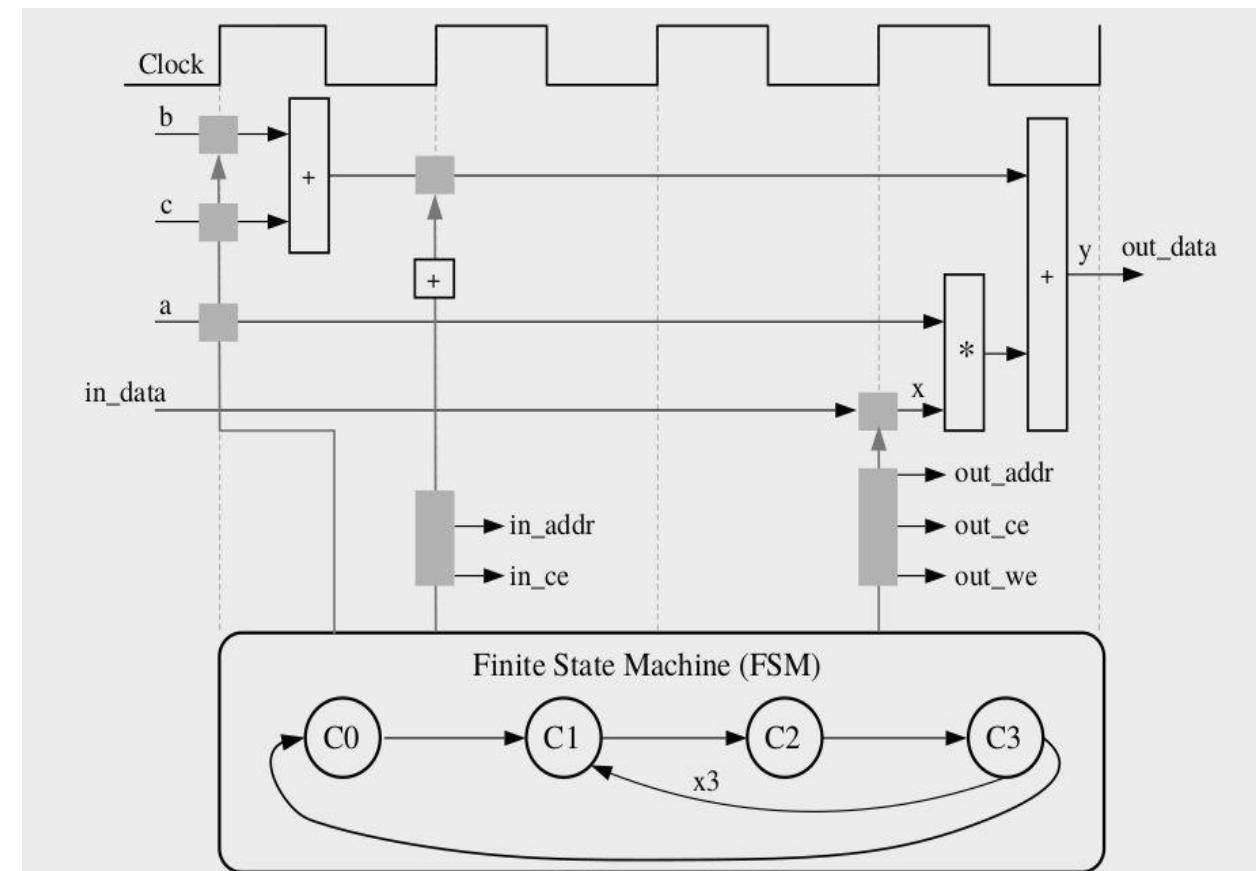
# Module 2: High Level Synthesis with Vivado-HLS

# High Level Synthesis (HLS)



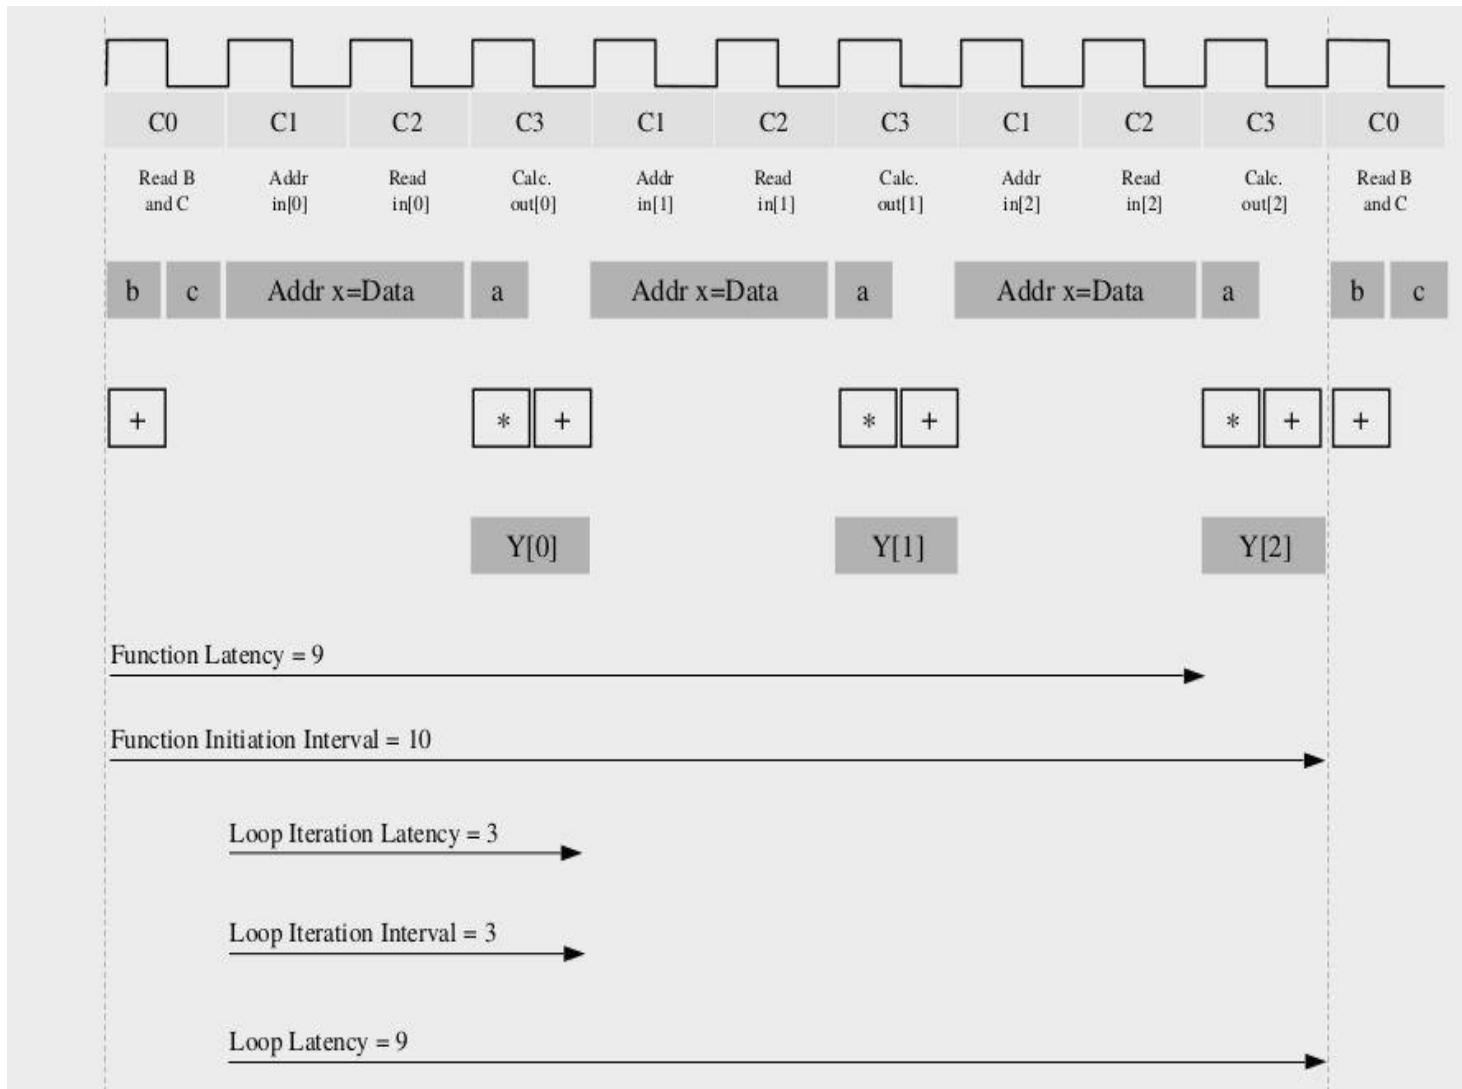
# Example

```
void foo(int in[3], char a, char b, char c, int out[3]) {
    int x,y;
    for (int i=0; i<3; ++i) {
        x = in[i];
        y = a*x + b + c;
        out[i] = y;
    }
}
```



Source: Xilinx Web

# Learning Goals?



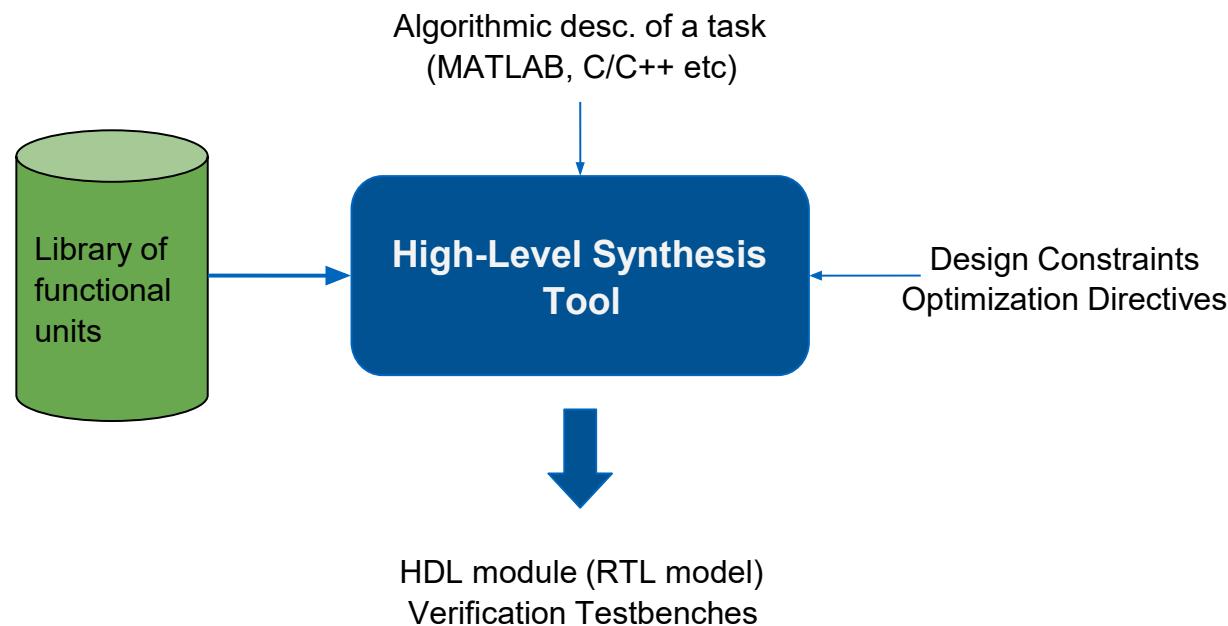
# HLS Benefits

- Design at higher abstraction
  - improved productivity for HW designers
    - employing Model-Based-Design principles
    - rapid system prototyping; shorter time-to-market
  - improved system performance for SW designers
    - HW/SW CO-design to accelerate applications
- Better Design Space Exploration
  - optimize resource, perf. trade-off early in design cycle

# HLS Limitations

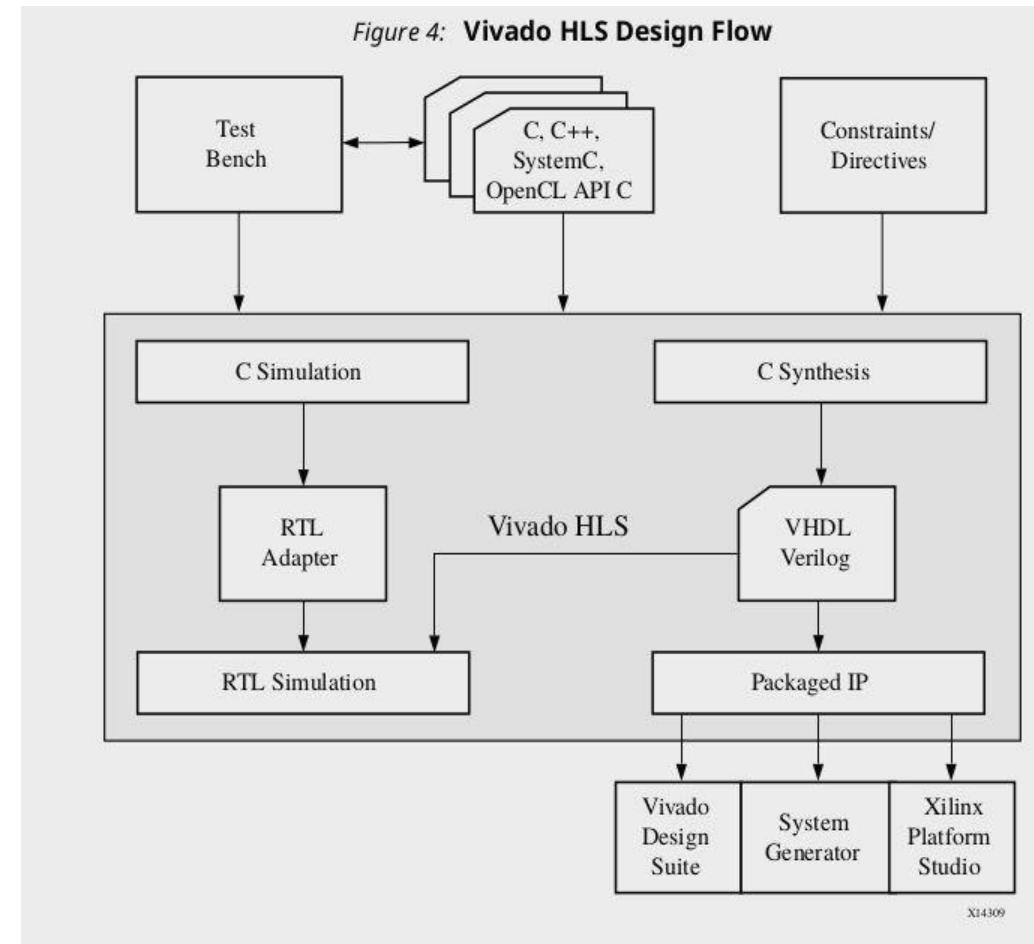
- Some SW concepts do not map well into HW
  - Datatypes: No fixed datatypes
  - Pointers: No heap for dynamic memory
  - Recursion: No notion of execution stack
  - Communication: No notion of shared memory
- Generated RTL hard to extend / modify
  - difficult to debug if RTL verification fails

# HLS abstract tool view



# Vivado HLS Flow

- Inputs:
  - C function
  - Constraints
  - Directives
  - C testbench
- Outputs:
  - RTL code in HDLs, SystemC etc.
  - RTL testbench
  - Report files



Source: Xilinx Web

X14309

# Vivado-HLS: Implementation Aspects (1)

- Top-level function args into
  - RTL I/O ports
  - supported Interfaces (AXI)
- C functions into blocks in RTL hierarchy
  - hierarchy of sub-funcs -> hierarchy of RTL modules
  - all instances of a function use same RTL block
- Loops into blocks in RTL hierarchy with control path
  - nested-loops -> hierarchy of RTL modules

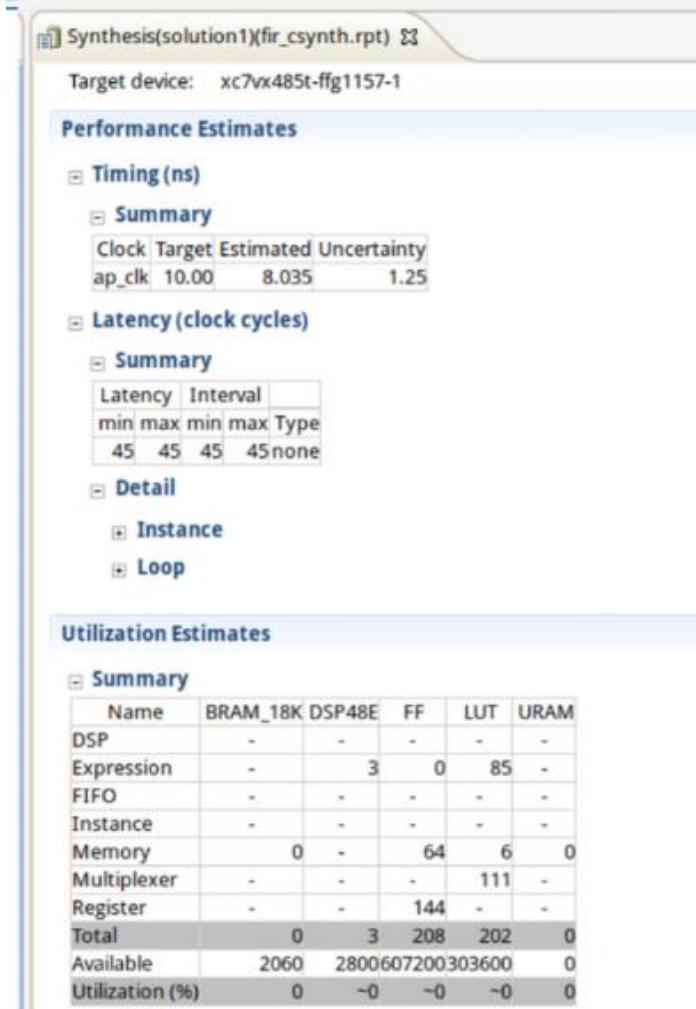
## Vivado-HLS: Implementation Aspects (2)

- Loops are kept rolled by default
  - RTL logic (as component) for one iteration of loop
  - all iterations scheduled in sequence using the same logic
- Arrays synthesize to Block-RAM by default
  - Arrays in I/O ports as external block-RAMs
  - FIFOs, HW Registers, Distributed RAM also possible

# Synthesis

## Synthesis reports:

- Estimated Timing
- Utilization
- Loop/Module:
  - iteration latency
  - initiation interval



The screenshot shows a synthesis report window with the following sections:

### Performance Estimates

- Timing (ns)**
  - Summary**

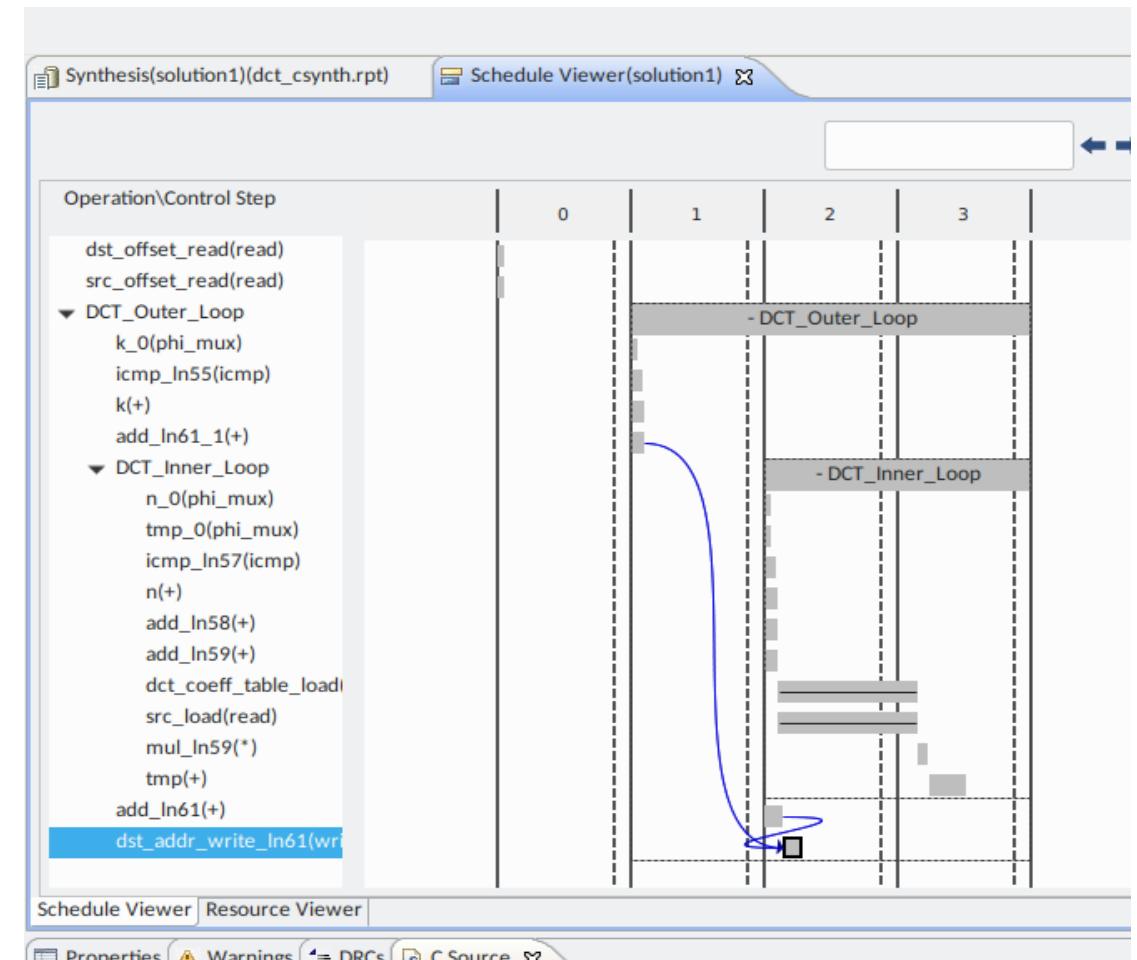
Clock	Target	Estimated	Uncertainty
ap_clk	10.00	8.035	1.25
  - Latency (clock cycles)**
    - Summary**

Latency	Interval
45	45 45 45
    - Detail**
      - Instance**
      - Loop**

# Analysis

## Reports:

- Data Dependencies within Modules
- Various Latencies such as:
  - Loop iteration latency
  - Loop initiation interval



- **Strategies:**
  - Instruct a task to execute pipelined
  - Specify latency for completion of funcs/loops
  - Specify limit on number of resources used
  - Select a specific I/O protocol for optimal integration
  - Efficient structuring of data items

# Directives

Type	
Kernel Optimization	<ul style="list-style-type: none"><li>• <code>pragma HLS allocation</code></li><li>• <code>pragma HLS expression_balance</code></li><li>• <code>pragma HLS latency</code></li><li>• <code>pragma HLS reset</code></li><li>• <code>pragma HLS resource</code></li><li>• <code>pragma HLS top</code></li></ul>
Function Inlining	<ul style="list-style-type: none"><li>• <code>pragma HLS inline</code></li><li>• <code>pragma HLS function_instantiate</code></li></ul>
Interface Synthesis	<ul style="list-style-type: none"><li>• <code>pragma HLS interface</code></li><li>• <code>pragma HLS protocol</code></li></ul>
Task-level Pipeline	<ul style="list-style-type: none"><li>• <code>pragma HLS dataflow</code></li><li>• <code>pragma HLS stream</code></li></ul>
Pipeline	<ul style="list-style-type: none"><li>• <code>pragma HLS pipeline</code></li><li>• <code>pragma HLS occurrence</code></li></ul>
Loop Unrolling	<ul style="list-style-type: none"><li>• <code>pragma HLS unroll</code></li><li>• <code>pragma HLS dependence</code></li></ul>
Loop Optimization	<ul style="list-style-type: none"><li>• <code>pragma HLS loop_flatten</code></li><li>• <code>pragma HLS loop_merge</code></li><li>• <code>pragma HLS loop_tripcount</code></li></ul>
Array Optimization	<ul style="list-style-type: none"><li>• <code>pragma HLS array_map</code></li><li>• <code>pragma HLS array_partition</code></li><li>• <code>pragma HLS array_reshape</code></li></ul>
Structure Packing	<ul style="list-style-type: none"><li>• <code>pragma HLS data_pack</code></li></ul>

[https://download.amd.com/docnav/documents/aem/xilinx2019\\_1-ug1253-sdx-pragma-reference.pdf?utm\\_source=chatgpt.com](https://download.amd.com/docnav/documents/aem/xilinx2019_1-ug1253-sdx-pragma-reference.pdf?utm_source=chatgpt.com)

## M2.1 Lab tasks

# Finite Impulse Response (Getting started)

- Intro to Vivado HLS tool
- Synthesize an initial solution for FIR filter (C func)
- Steps:
  - Set up an HLS project
  - Source code for FIR filter in C
  - Validate the C source
  - Create and synthesize a solution
  - Analyze the synthesized HW
  - Package the RTL as IP

## HLS of a DCT Filter

- Discrete Cosine Transform (DCT) is heavily used in signal-processing applications
- Steps:
  - Start off with an initial C-level implementation
  - Analyze the performance bottlenecks
  - Optimize via directives to improve performance
  - Package the final generated IP

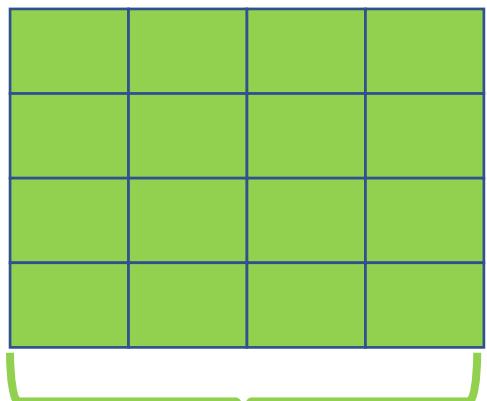
# Module 3: Acceleration of embedded AES Cryptography with HLS

# M3.1 Background

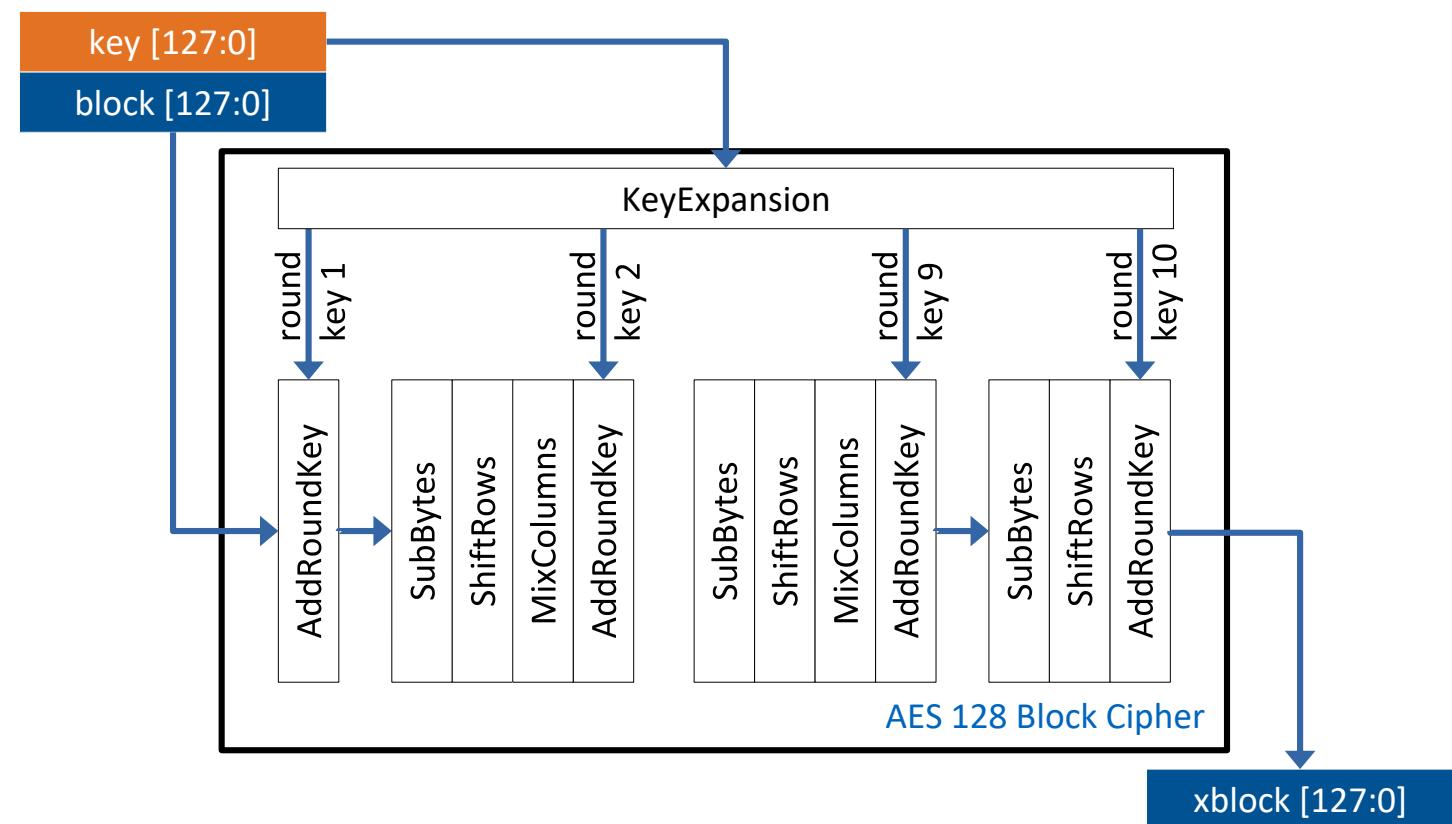
# Advanced Encryption Standard (AES)

- Essentially a symmetric data-processing algorithm
  - Block Cipher AES-128 (128-bit blocks)  
→ requires add. Block Mode of Operation (BCMO) to en-/decrypt messages of multiples blocks
- In general compute-intensive
  - High optimization potential for HLS
    - Regular data-flow oriented nature
    - Symmetry allows usage of same hardware for both directions
    - **Please don't do this in HLS in the real life! -> Side-Channels**

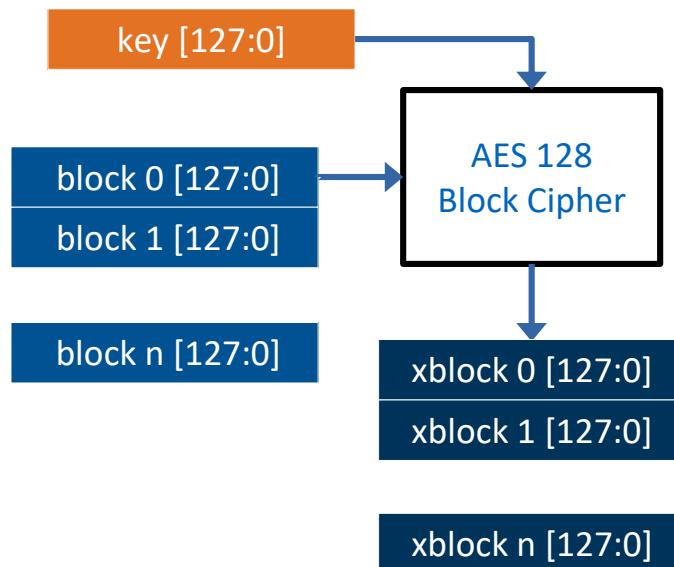
# AES-128



- 128-bit IO and key. Input mapped to **4\*4 byte state**
  - **KeyExpansion**: Generate individual key for each round
  - **AddRoundKey**: state is XORed with round key
  - **SubBytes**: Each state byte's value is substituted by a mapped value (can be a LUT or calculated) (non-linear)
  - **ShiftRows**: Shifting rows of state matrix
  - **MixColumns**: Mixing operation on columns of state matrix
- > computationally heavy



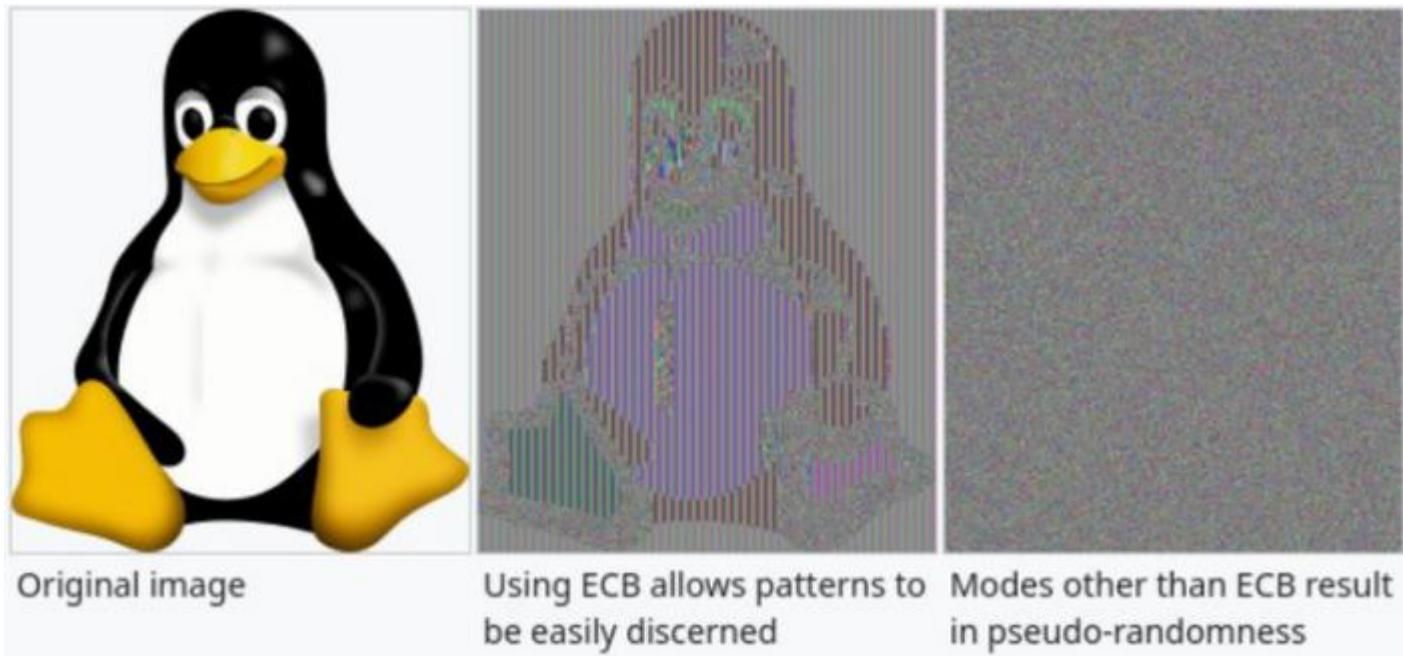
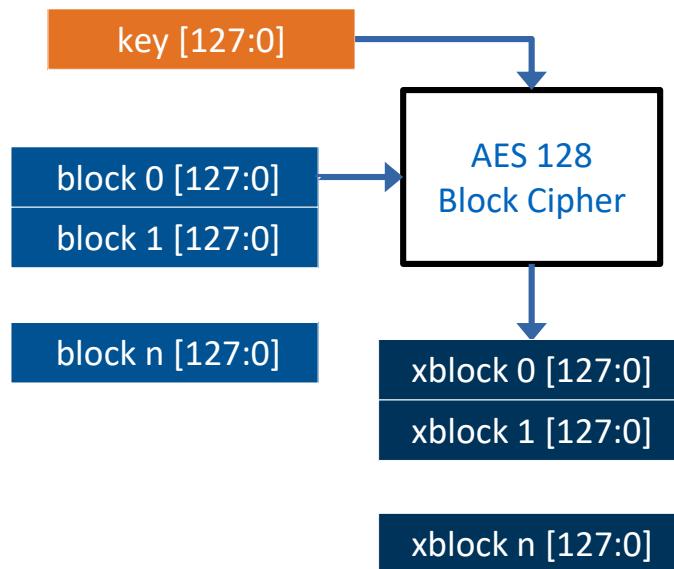
# Encrypting a message: Straight-forward



- Consider message  $m$  with  $n$  128-bit blocks of information of which some bear the same information, i.e. their 128-bit value is equal
- Block-by-block transformation with AES128 would yield encrypted but recurring encrypted blocks.

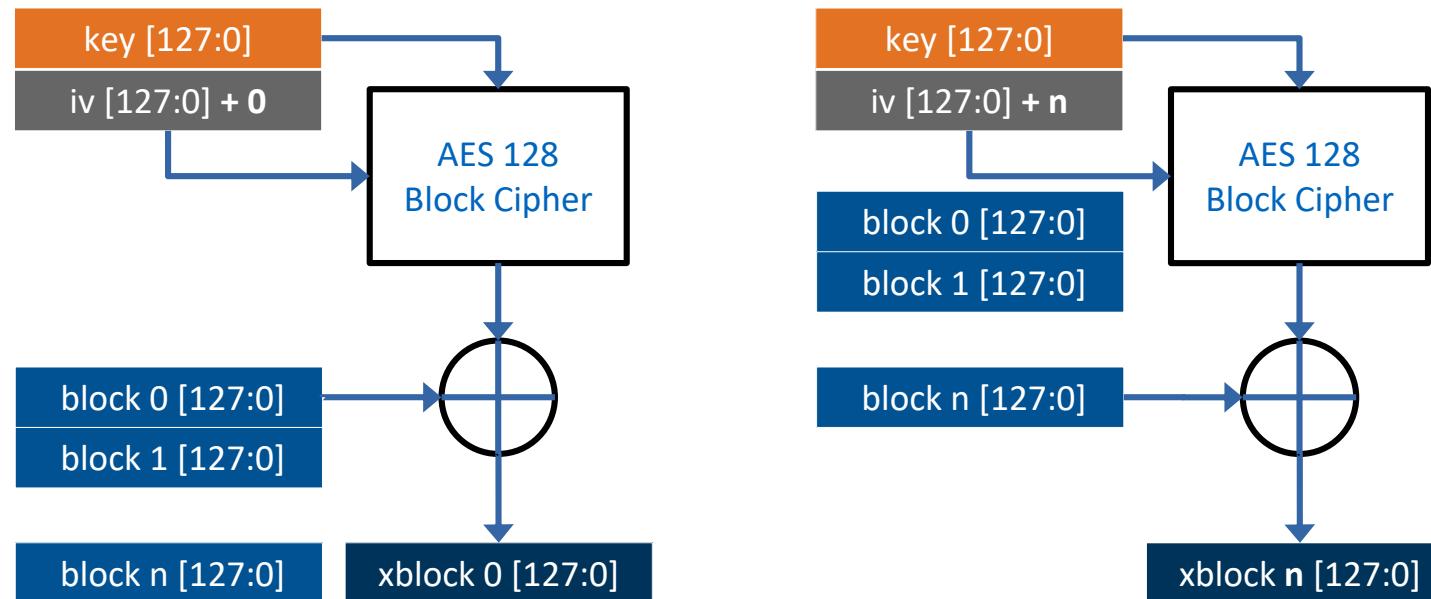
→ **Very bad!**

# Encrypting a message: Straight-forward



Comparison between original, ECB and otherwise encrypted messages  
[https://upload.wikimedia.org/wikipedia/commons/thumb/9/96/Tux\\_encrypted\\_ecb.png/196px-Tux\\_encrypted\\_ecb.png](https://upload.wikimedia.org/wikipedia/commons/thumb/9/96/Tux_encrypted_ecb.png/196px-Tux_encrypted_ecb.png)

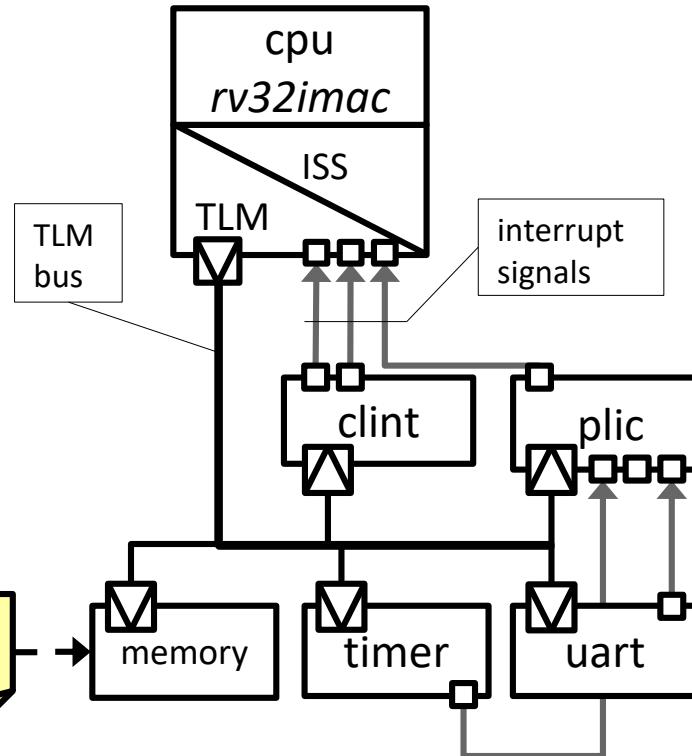
# Encrypting a message: BCMO



Block Cipher Mode of Operation (BCMO) enables Block Ciphers (like AES128) to be used on multi-block messages. Here: Counter (CTR) + Initiation Vector (IV)

## **M3.2 Task 1: AES-SW**

# Task 1: Pure Software implementation



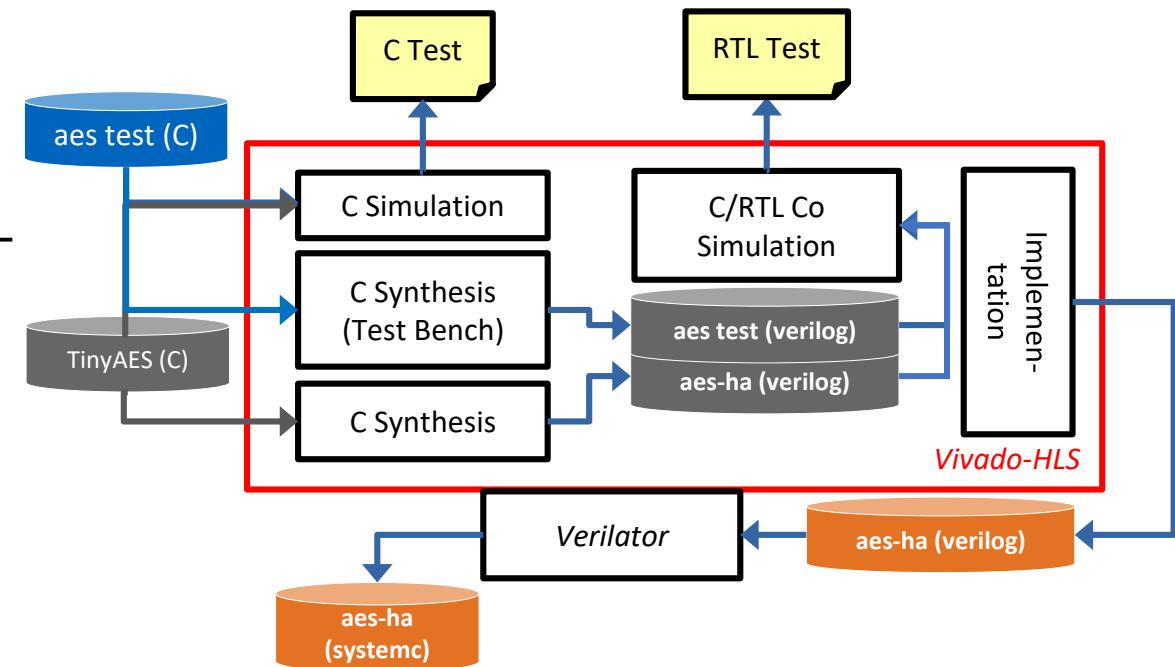
- Decrypting a message  $m$ , containing a secret ASCII-based picture.
- $m$  can be resembled by a  $46 \times 192$  byte =  $46 \times (12 \times 128)$  bit matrix.
- $m$ , private key (AES key) and public key (IV) are provided

Task: Application Profiling the AES software library as a software implementation on EDAduino. **#cycles spent for this solution?**

## M3.2 Task 2: AES-HA

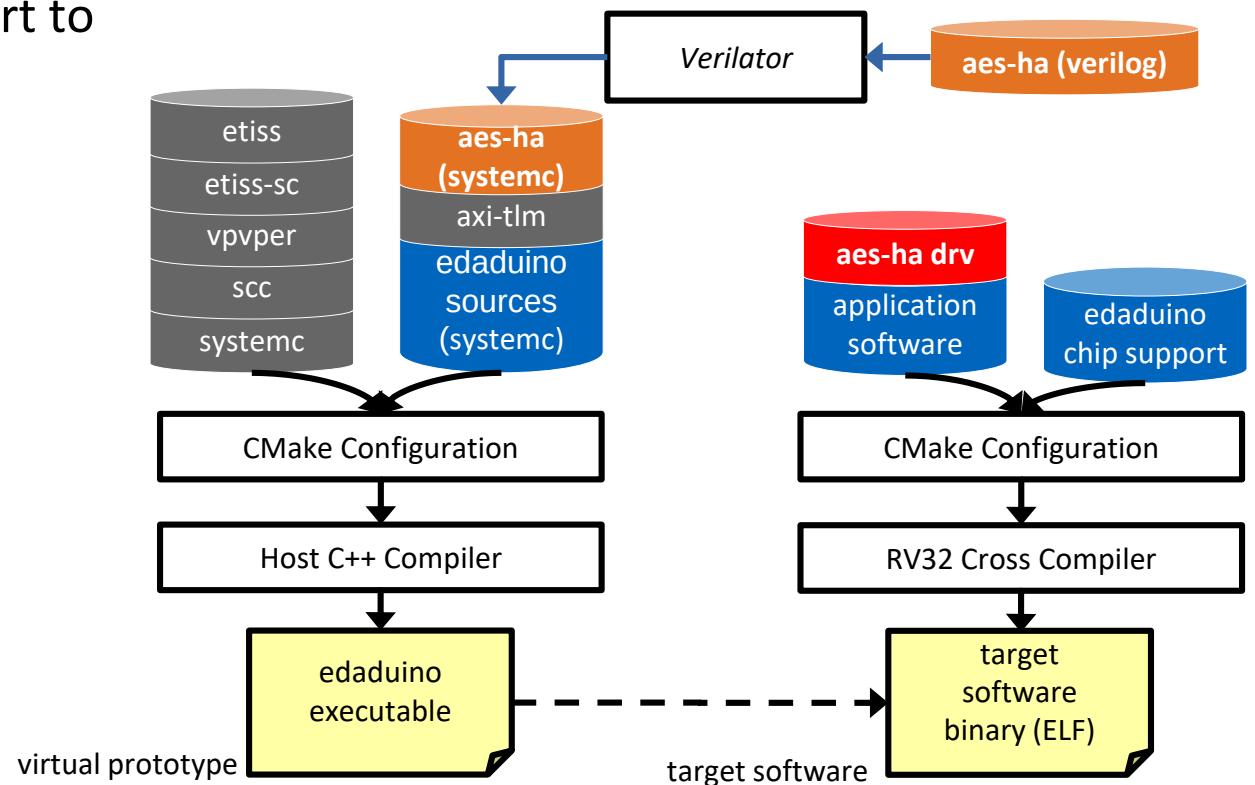
# Extended Design Flow: AES-CTR to AES-HA

- 1) AES-CTR (AES128 w/ CTR BCMO) as C software library
- 2) Vivado-HLS builds RTL (normally we would export to a Xilinx FPGA design suite)
- 3) Verilator builds cycle-accurate SystemC model from RTL  
-> Tests at each design stage

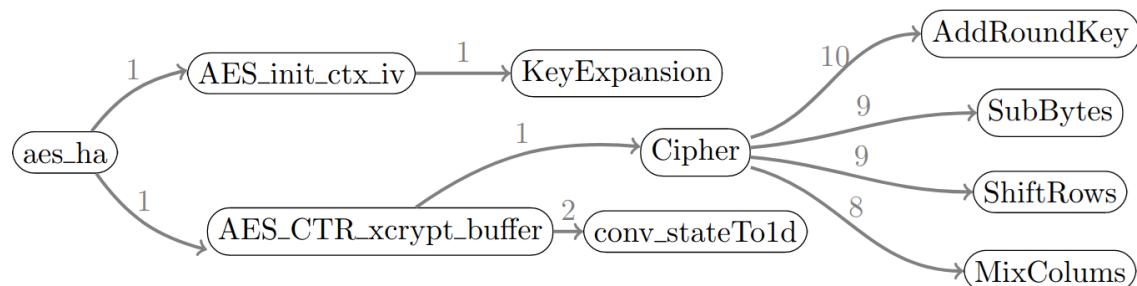


# Extended Design Flow: AES-CTR to AES-HA

- 1) AES-CTR (AES128 w/ CTR BCMO) as C software library
  - 2) Vivado-HLS builds RTL (normally we would export to a Xilinx FPGA design suite)
  - 3) Verilator builds cycle-accurate SystemC models from RTL
- > Tests at each design stage



## Task 2: HLS AES-SW to AES-HA



Task: Generate and optimize a hardware accelerated solution with Vivado-HLS.

### Rules:

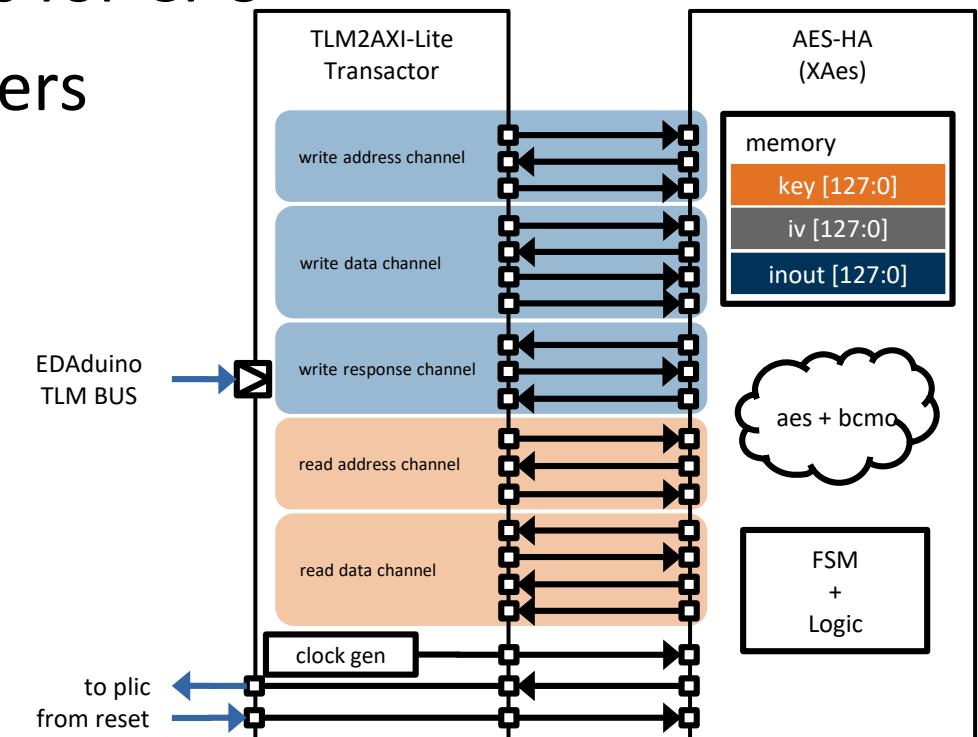
- 1) Do not change nor introduce additional directives to the top function `aes_ha.c`,
- 2) the design is implemented for the Xilinx FPGA chip `xc7s15-ftgb196-2`,
- 3) no resource usage shall exceed 100%,
- 4) a maximum latency of 500 clock cycles, and
- 5) the solution must pass the Vivado-HLS C/RTL Cosimulation test)

## M3.3 Task 3: AES-HW Integration

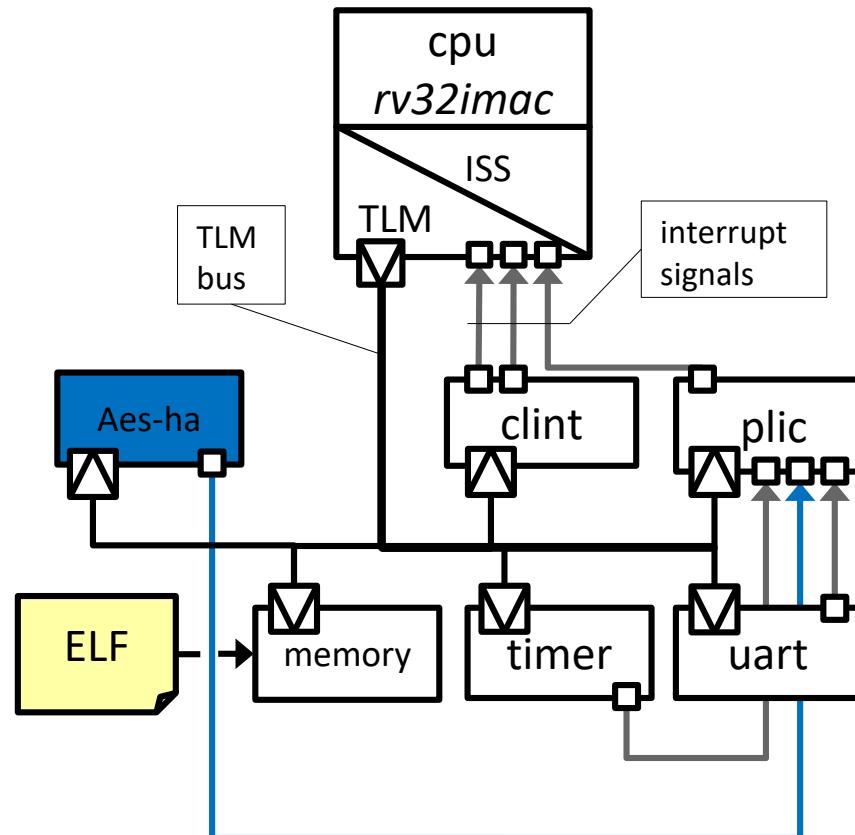
# EDAduino AES-HA Integration (Basic)

Transactor provides memory-mapped registers for CPU and routes data flow AES-HA (VRTL) and registers

- Key, IV, readable and writeable by CPU
- IV updated by CTR after each AES-HA run
- Interrupt gen. by FSM signals to CPU



# Task 3: HA No DMA Implementation



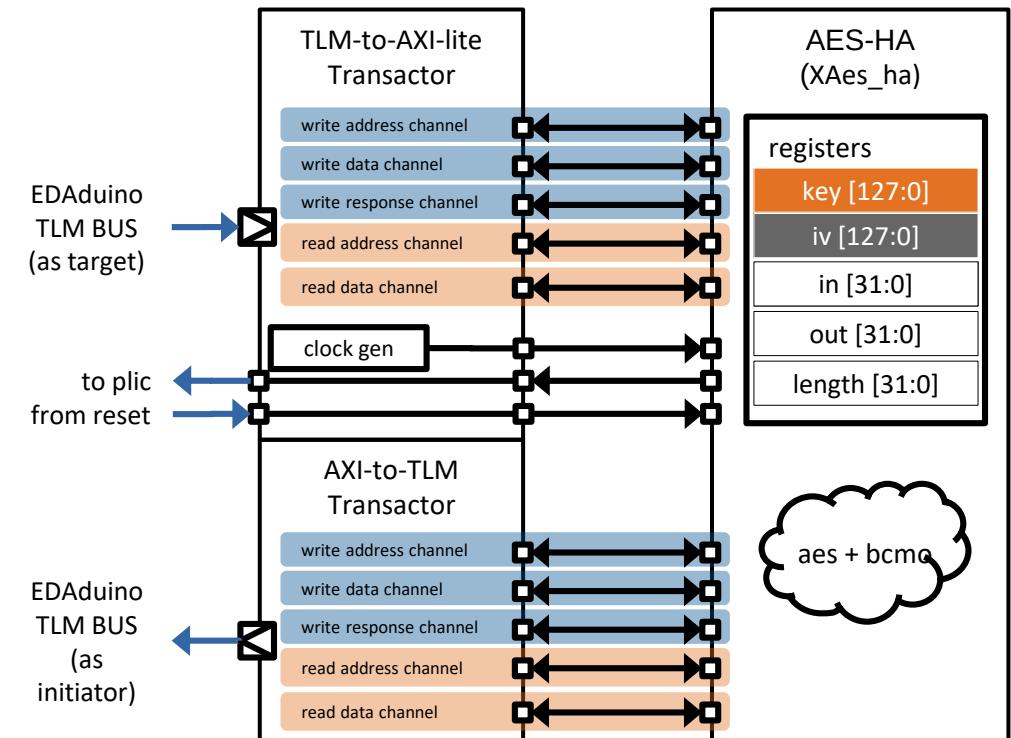
- No Direct Memory Access from **aes-ha**
- AXI-lite target port for control
- Same AXI-lite target port for I/O
- Driver library generated by Vivado-HLS
- CPU has to ...
  - read message from memory
  - write input register
  - issue start command
  - synchronize (wait)
  - read output register
  - write to memory

Task: Integrate, drive and evaluate the hardware accelerated cryptography on EDAduino. **#cycles spent for this solution?**

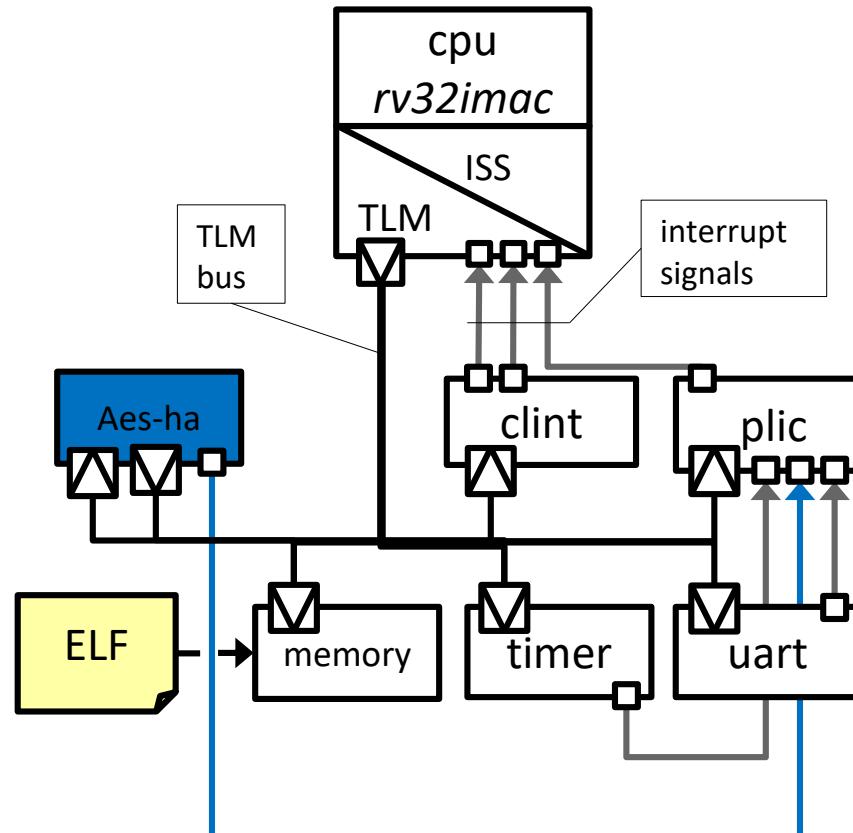
# EDAduino AES-HA Integration (DMA)

AES-HA with AXI4-M (Initiator) port for direct memory access of AES-HA

- Key, IV, readable and writeable by CPU
- Message length, Input address and output address writeable
- IV updated by CTR after each AES-HA run
- Interrupt gen. by FSM signals to CPU
- No CPU memory operations for I/O



# Task 3: HA with DMA Implementation



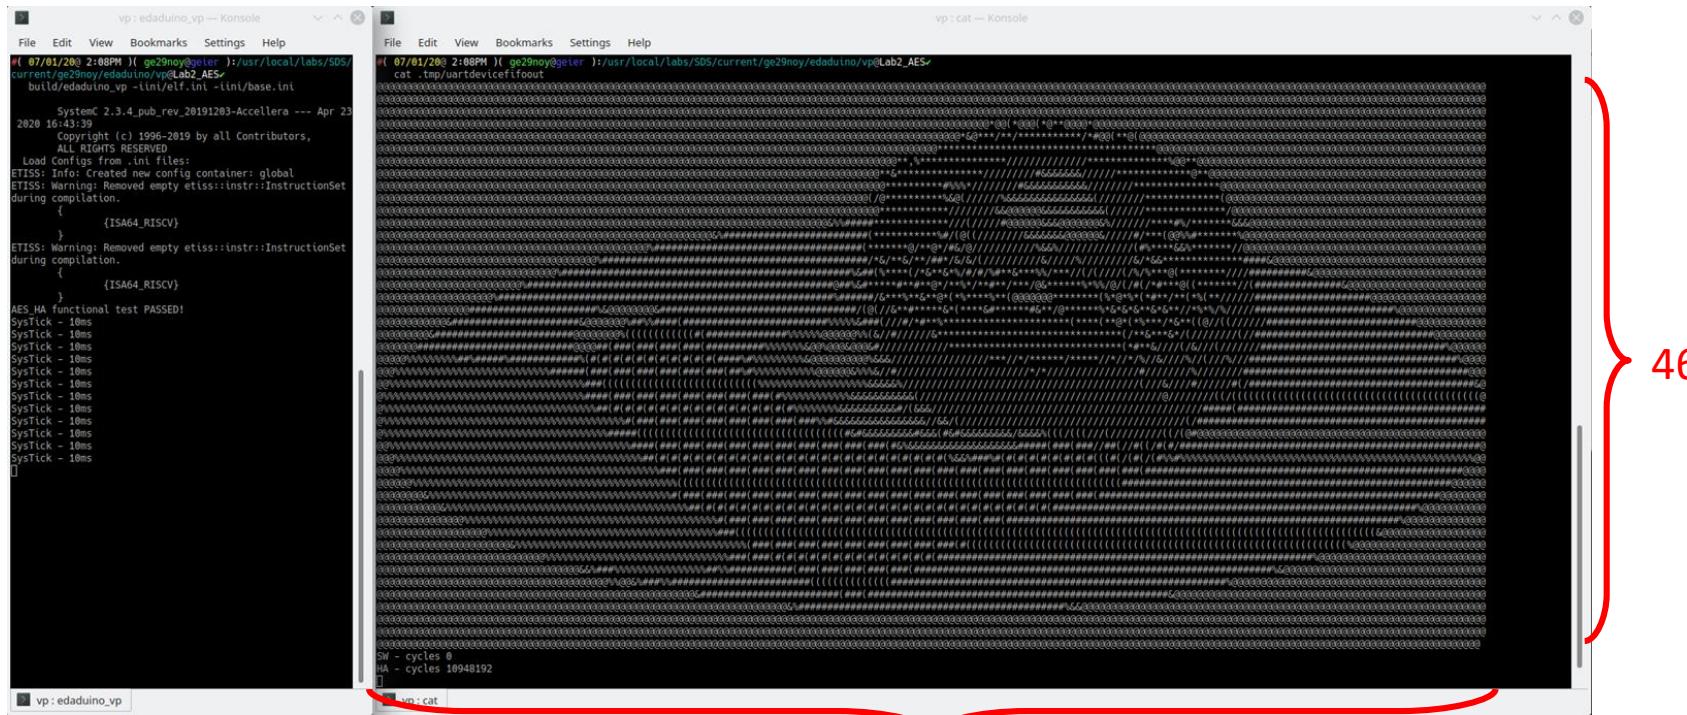
- Direct Memory Access from **aes-ha**
- AXI-lite target port for control
- AXI-M initiator port for I/O
- Driver library generated by Vivado-HLS

- CPU has to ...

- write message address to input register
- write output destination address
- issue start command
- synchronize (wait)

Task: Integrate, drive and evaluate the hardware accelerated cryptography on EDAduino. **#cycles spent for this solution?**

# Reference: Decrypted Message



192

4

## Exam Preparation

- Solution Slides provided for the 3 modules
- Knowledge of Vivado HLS tool and most common directives
- Exam from last WS24 in HLS to practice (in a zip folder on TUWEL )  
(similar expectations in 90mins)

## INFO

The registration for the second lab exam is already open!

# Working on your own

# Options for Accessing the Lab Material

- **Lab Access:**
  - TILab accounts are generated by now and should be available next week (Room 4)
  - Access with your student card (<https://www.tilab.tuwien.ac.at/howto.shtml>)
- **Virtual Machine:**
  - Using the TILab account, a virtual machine containing all sources and dependencies can be downloaded
- **Source Files on gitlab:**
  - A TUWEL assignment was opened asking for gitlab IDs. We will use this to give you access to the material
- **Lab Manual will be uploaded to TUWEL**

## Warning

Warning: This lab has been reworked  
recently. So you might encounter Bugs.  
Please report them!!!

## 2nd Warning

Warning: The VM has been recently created and is giant (119GB  
uncompressed) and 52 GB compressed!

If you want to use it, test it ASAP!

Decompressing will take ages if you don't find a tool with xz algorithm  
(USE LINUX ^^)

**Thank you for your attention!**