

Component description

Objective

This component aims at encapsulating some code written in C into a hardware component, called *virtual coprocessor wrapper*. The objective targetted is to have an estimation of the performances when a task is executed in hardware, without having to write a real coprocessor.

Modelisation

The coprocessor is modeled by a hardware component, the *wrapper*, which contains registers, a transition function and a Moore generation function. It can be interfaced with a `vci_mwmr_controller` that allows access to any number of MWMR channels in memory. The C function corresponding to the hardware task is associated to a posix thread, which is executed in parallel with the simulator for the architecture. This task thread is launched by the *wrapper* on reset.

The task thread communicates with the *wrapper* with two fifo channels:

- the `cmd` channel, which transmits to the *wrapper* the SRL commands:
 - ◆ `srl_mwmr_read(channel, buffer, size)`
 - ◆ `srl_mwmr_write(channel, buffer, size)`
 - ◆ `srl_busy_cycles(n_cycles)`
- the `rsp` channel, which is used to transmit the response to the read requests



To avoid race conditions, a single lock is used for both fifos `cmd` and `rsp`.

The *wrapper* is implemented as a four-state automaton, whose role is to execute the SRL commands transmitted by the hardware task.



1. In IDLE state, the automaton tests if a command is available in the `cmd` fifo.
 - ◆ If yes, it gets the parameters and goes in a state corresponding to the type of the command to execute it.
 - ◆ If no, it waits until a command is available
2. In READ state, the automaton transfers data from the MWMR controller directly into the buffer of the task thread (1 word/cycle). When the last word has been transfered, the automatod signals the completion to wake up the task thread, and returns in IDLE state.
3. In WRITE state, the automaton makes the transfer from the task thread to the MWMR controller (1 word/cycle), and returns in idle when the transfer is complete.
4. In BUSY state, the automaton decreases a counter at each cycle until the counter reaches 0.

Synchronization

The *wrapper* behaves as if all the SRL commands were blocking for the task thread, which is well adapted to sequential coprocessors, which never execute several commands in parallel:



However, to avoid potential useless switches between the threads, only the read commands really are blocking. Thus, if a task thread only makes writes, it will continue executing until the `cmd` fifo is considered full (the default value is 50 commands, but it could be anything).

The two threads are synchronized with two `pthread_cond_t` conditions: `task_cond` and `wrapper_cond`. Here is the simplified execution scheme on one (host) processor, for a task thread making only reads:

- On creation, the task thread is executed
- The task thread makes a read: it sends a command and then try to read a response; as there is none, it waits on `task_cond`
- The wrapper thread is executed, receives the read request, and process it. When the response is ready, it signals `task_cond`
- The wrapper thread tries to read another command. As there is none, it waits on `wrapper_cond`
- The task thread is executed, and makes another read, etc.

Here is the simplified execution scheme on one (host) processor, for a task thread making only writes:

- On creation, the task thread is executed
- The task thread makes a write: it sends a command and then continue its execution
- After a certain number of writes, the command fifo reaches its maximum value. The task thread signals `wrapper_cond` and waits on `task_cond`
- The wrapper thread is executed, and start processing all the writes commands.
- When there is no more commands in the `cmd` fifo, the wrapper thread signals `task_cond` and waits on `wrapper_cond`
- The task thread is executed, and makes another write, etc.

With several cores, the difference is that after the creation of the task thread, the wrapper thread can continue its execution, and theoretically consume the writes as they are produced, thus none of the threads never waits (though in practice the task thread will always be faster).

(A finir)

Component usage

Object

```
soclib::common::ProcessWrapper *process;
```

Instanciation

```
soclib::common::ProcessWrapper(  
    const std::string &cmd,  
    const std::vector<std::string> &argv );
```

`cmd`

The command to launch, it will be looked for in \$PATH

argv

The argv in a standart unix style. argv[0] should look like cmd.

Usage

```
std::vector<std::string> argv;
argv.push_back("date");
argv.push_back("+%s");
soclib::common::ProcessWrapper process("date", argv);

char buffer[128];
process.read( buffer, 128 );
std::cout << "Date is " << buffer << " seconds since EPOCH" << std::endl;
```