

# Core C++ 2021



## Design Patterns for Hardware Packet Processing on FPGAs

Haggai Eran

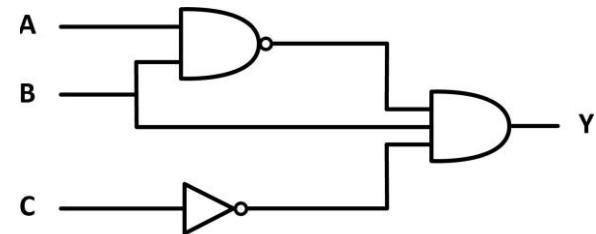
# TL;DR

- C++ for FPGA hardware design
- Packet processing example
- Design patterns and building blocks library

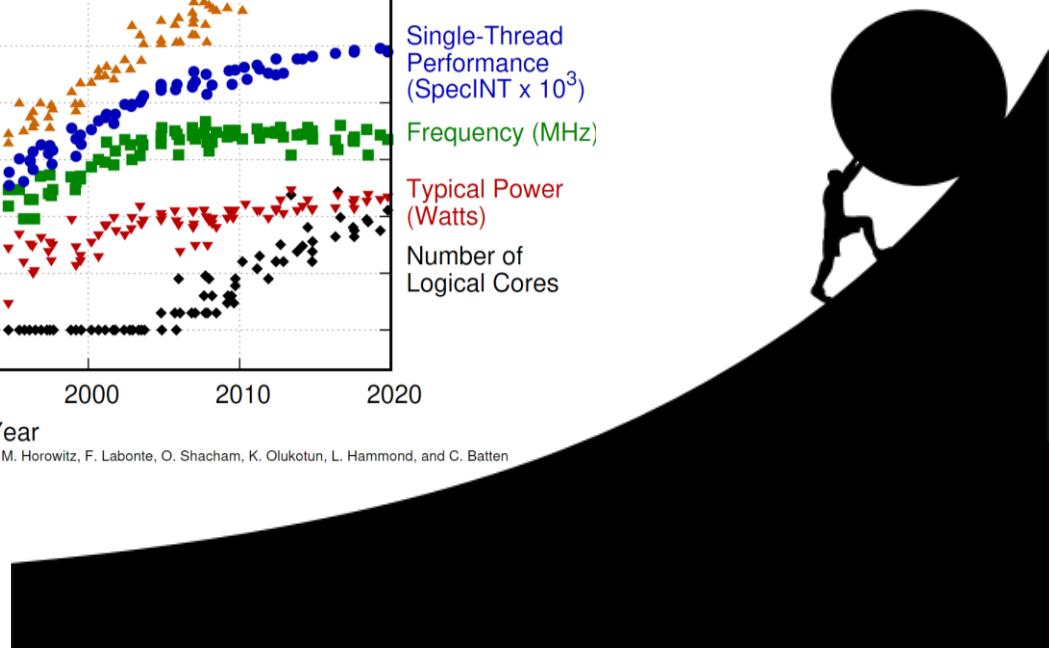
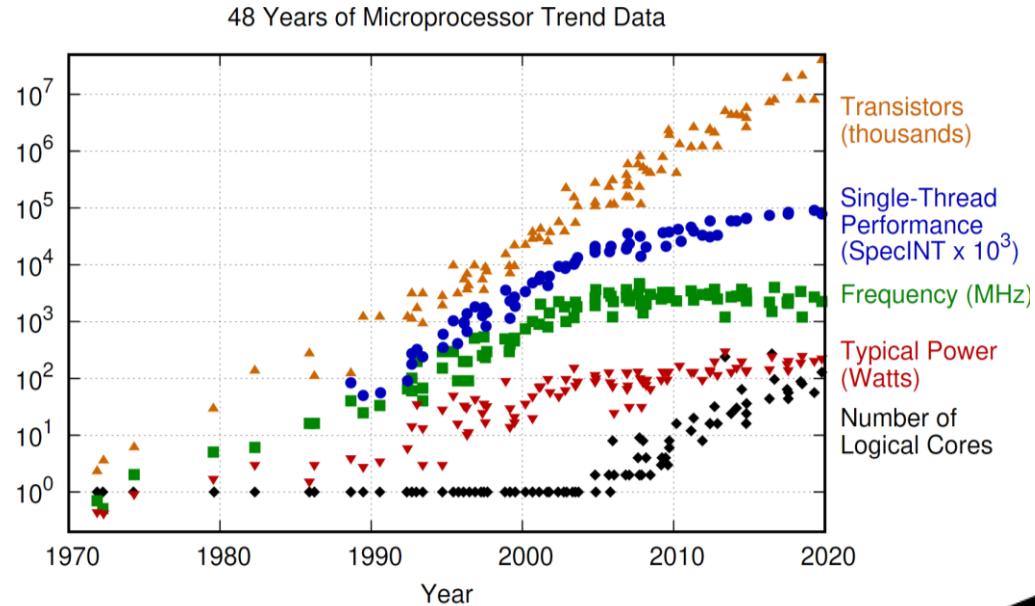
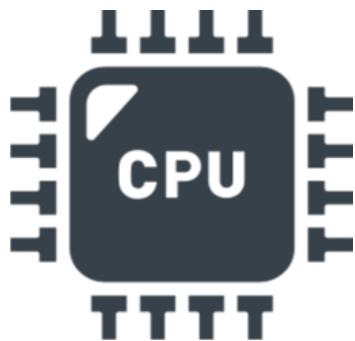
H. Eran, L. Zeno, Z. István and M. Silberstein,

"[Design Patterns for Code Reuse in HLS Packet Processing Pipelines](#),"

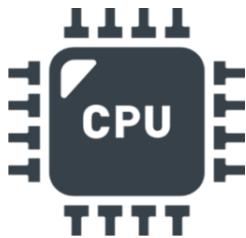
*2019 IEEE 27th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM), 2019*



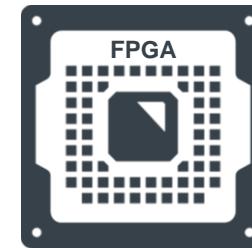
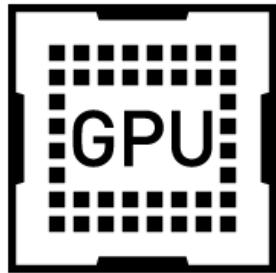
# Motivation: end of Dennard Scaling



# Accelerators & heterogenous computing



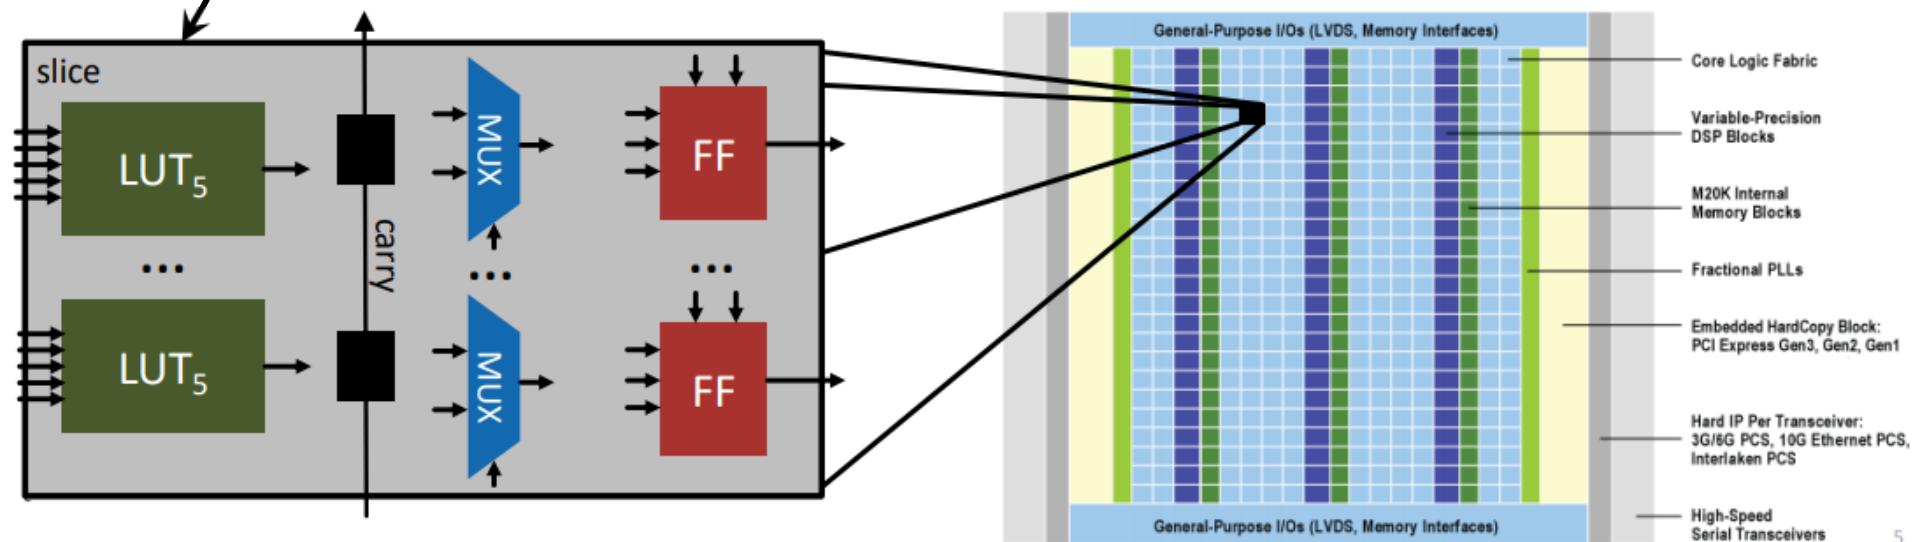
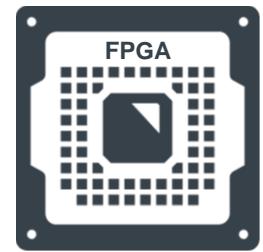
Cloud TPU v3



# What are FPGAs?



Input	Value
00000	0
00001	0
...	...
11111	1



# Network packet processing & FPGAs

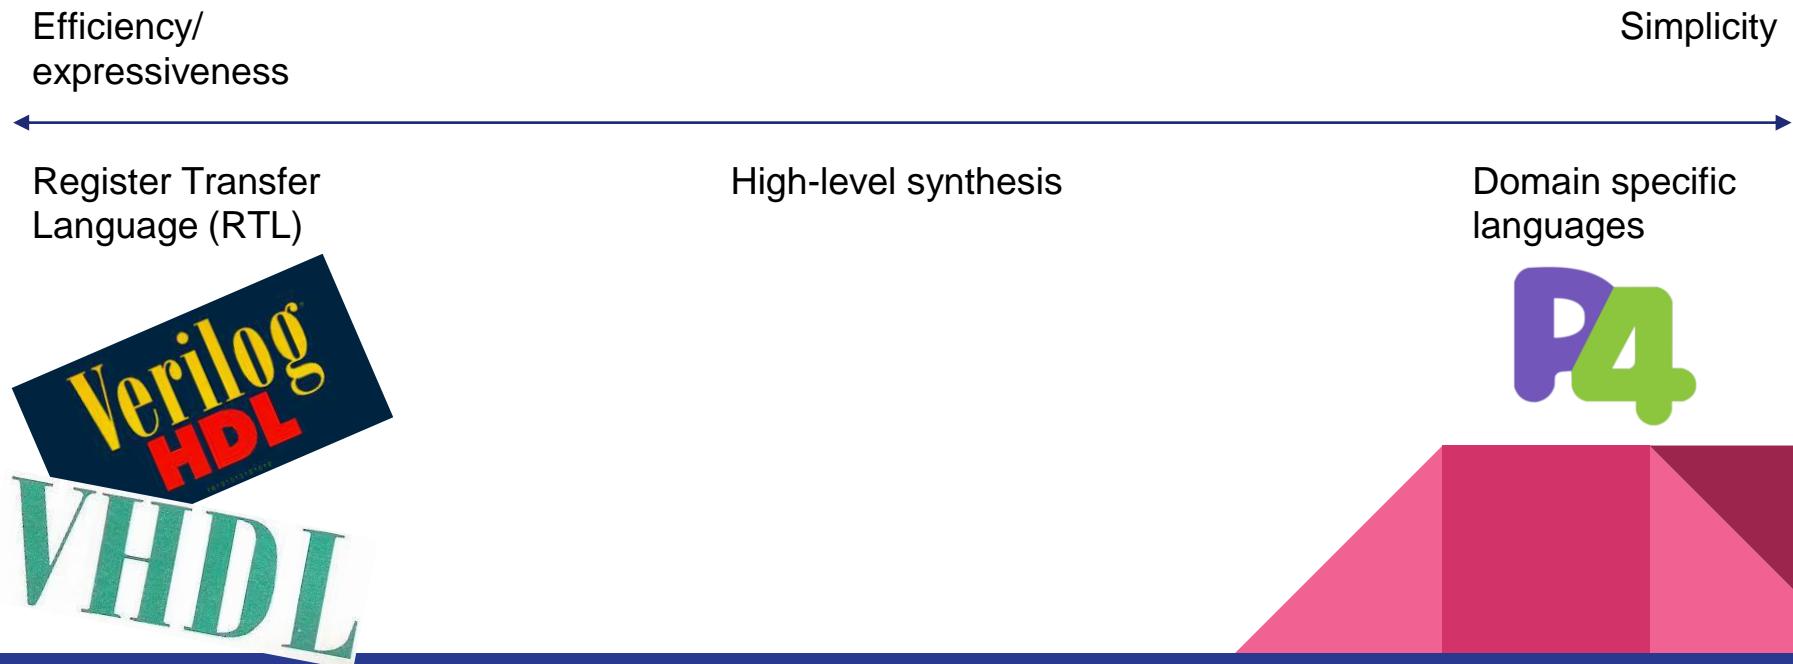
- High-throughput
- Low latency
- Predictability
- Flexibility

E.g.

- AccelNet on Microsoft Azure  
[Firestone et al. NSDI'18]



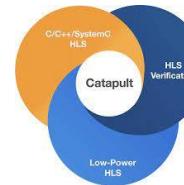
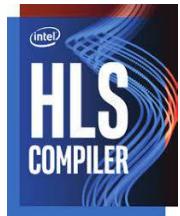
# Network packet processing on FPGAs programming alternatives



# High-level synthesis (HLS)



- Abstract underlying hardware details
  - Reuse a design on different hardware
- Rapid simulations
- Design space exploration



**Mentor®**  
A Siemens Business

**synopsys®**  
Symphony C

**cadence**

Stratus

# Why is it hard to build an HLS networking lib?

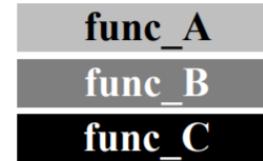
-  Only a subset of C++ is synthesizable.
  - No virtual functions
  - No dynamic memory allocation
  - No casting pointers
-  Vivado HLS 2018.2 supports C++11
-  Strict interfaces and patterns for performance.

# In this talk

- Vivado HLS background
- Legacy HLS – how others have used HLS for high-performance networking, limiting reuse
- ntl – template library for HLS packet processing and methodology used to develop it
- Quantitative comparison

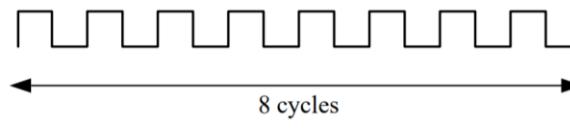
# Vivado HLS's dataflow optimization

```
void top (a,b,c,d) {  
    ...  
    func_A(a,b,i1);  
    func_B(c,i1,i2);  
    func_C(i2,d)  
  
    return d;  
}
```



# Vivado HLS's dataflow optimization

```
void top (a,b,c,d) {  
    ...  
    func_A(a,b,i1);  
    func_B(c,i1,i2);  
    func_C(i2,d)  
    ...  
    return d;  
}
```



func\_A | func\_B | func\_C

8 cycles

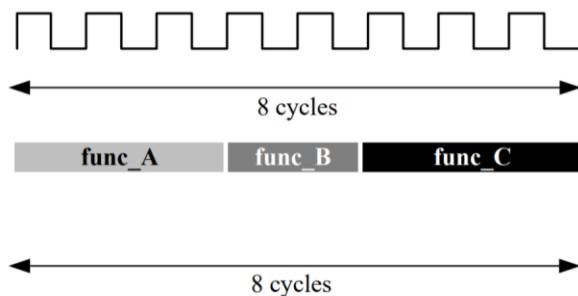
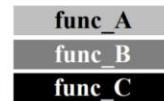
(A) Without Dataflow Pipelining

# Vivado HLS's dataflow optimization

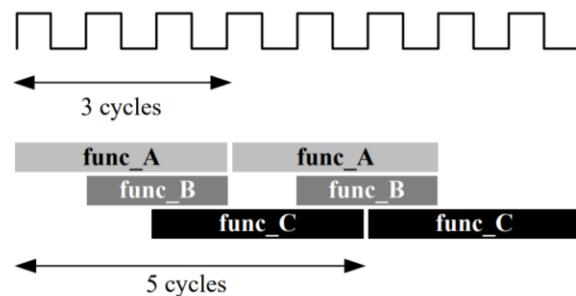
Figure 64: Dataflow Optimization

#pragma HLS dataflow

```
void top (a,b,c,d) {  
    ...  
    func_A(a,b,i1);  
    func_B(c,i1,i2);  
    func_C(i2,d)  
  
    return d;  
}
```



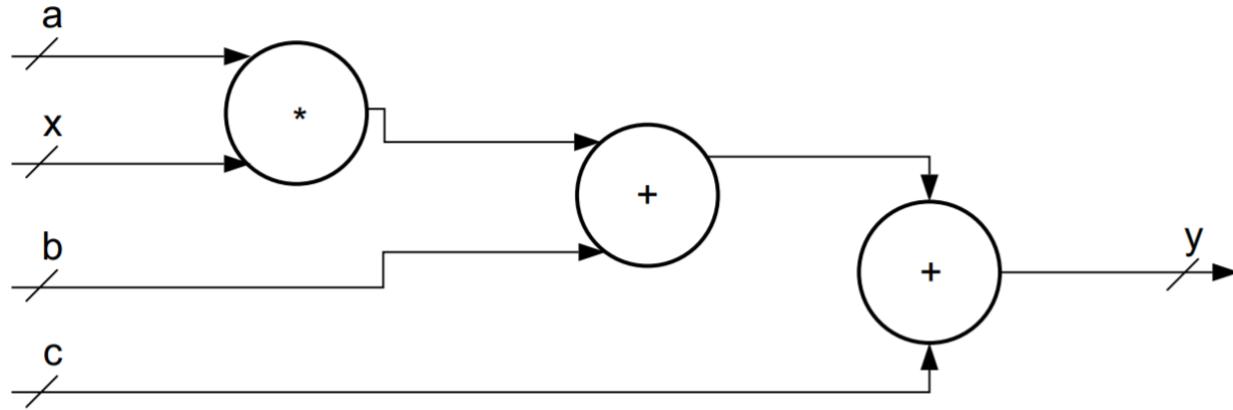
(A) Without Dataflow Pipelining



(B) With Dataflow Pipelining

# Vivado HLS's pipeline optimization

$$y = (a \times x) + b + c$$



# Vivado HLS's pipeline optimization

$$y = (a \times x) + b + c$$

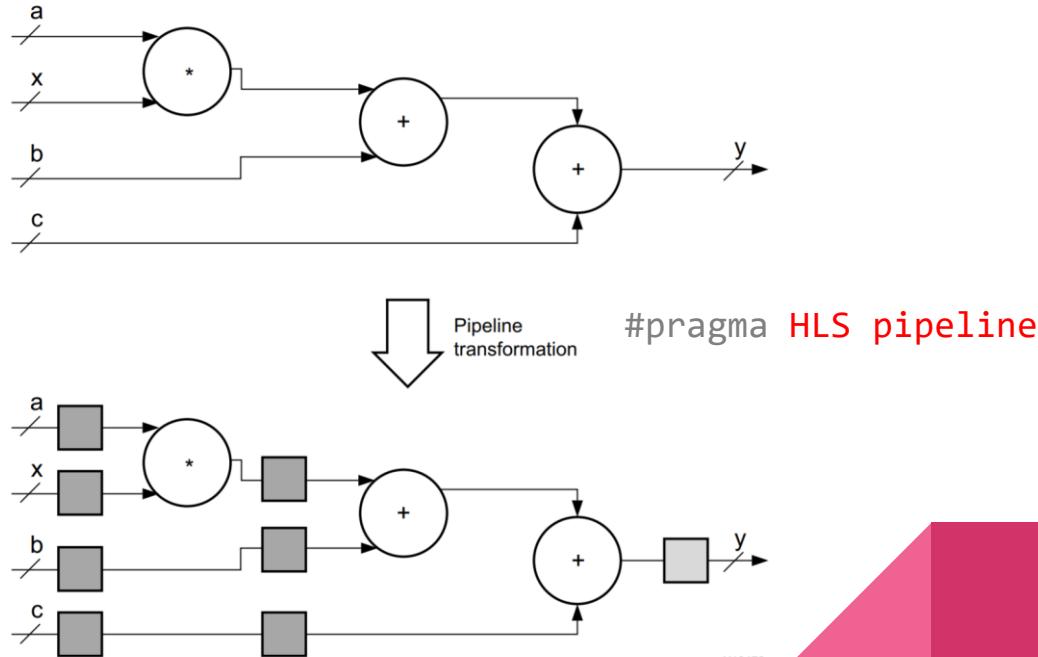


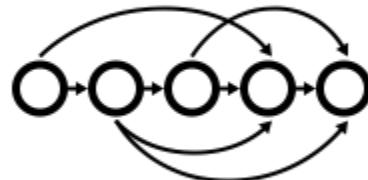
Figure 2-9: FPGA Implementation of a Compute Function

# Legacy HLS methodology

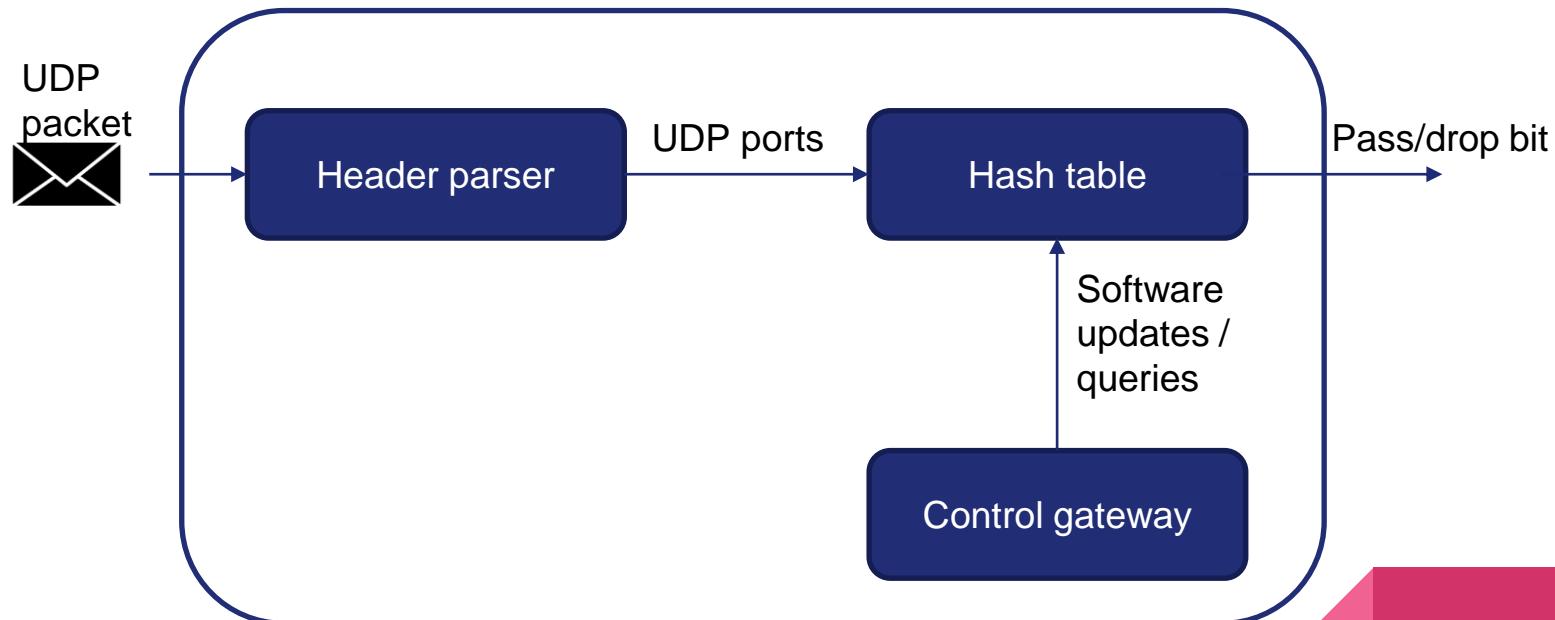
# Legacy HLS - how is HLS used for packet processing?

Data-flow design

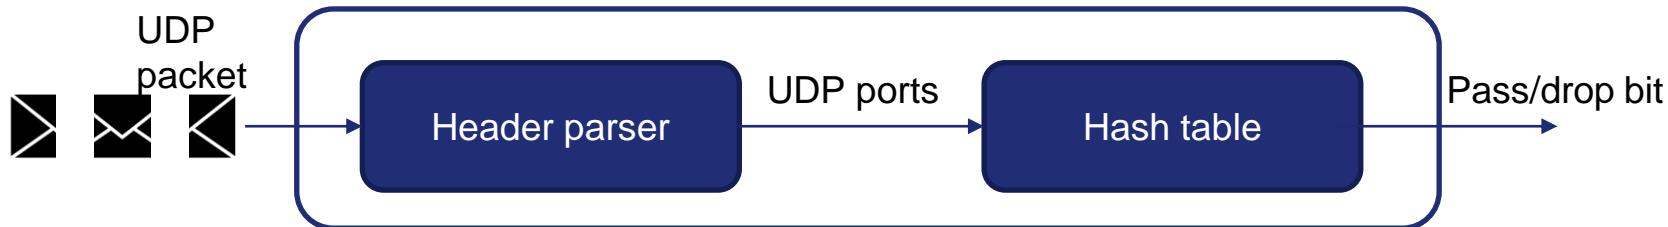
- A fixed graph of independent elements
- Operate on data when inputs are ready
- Examples: [[Blott '13](#)], [[XAPP1209 '14](#)], [[Sidler '15](#)], [ClickNP](#) [[Li '16](#)].



# Running example: UDP stateless firewall



# Packet interface: flits



# Legacy HLS: data-flow in Vivado HLS

```
struct axi_data {  
    ap_uint<256> data;  
    ap_uint<32> keep;  
    ap_uint<1> last;  
};
```

Describe arbitrary precision integers.

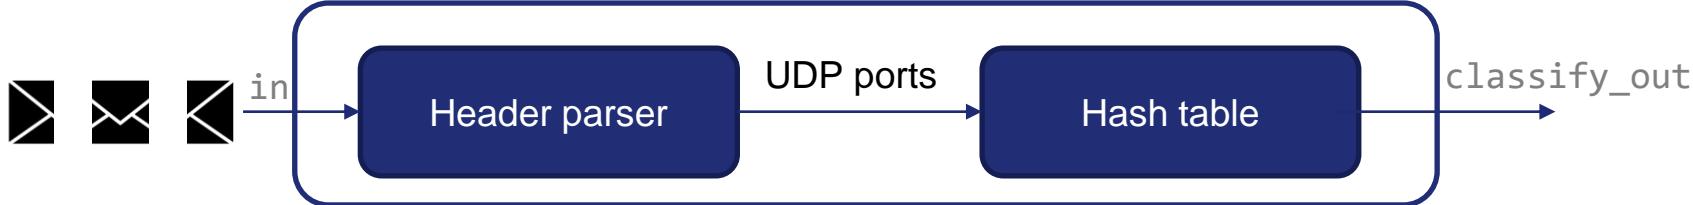
```
typedef hls::stream<axi_data> axi_data_stream;  
typedef hls::stream<ap_uint<1>> bool_stream;
```

Synthesized as FIFO.

# Legacy HLS: data-flow in Vivado HLS

```
void firewall_top(axi_data_stream& in, bool_stream& classify_out);
```

Passing by reference  
synthesizes a bus interface;  
direction is inferred



# Legacy HLS: data-flow in Vivado HLS

```
void firewall_top(axi_data_stream& in, bool_stream& classify_out)
{
    #pragma HLS dataflow
```

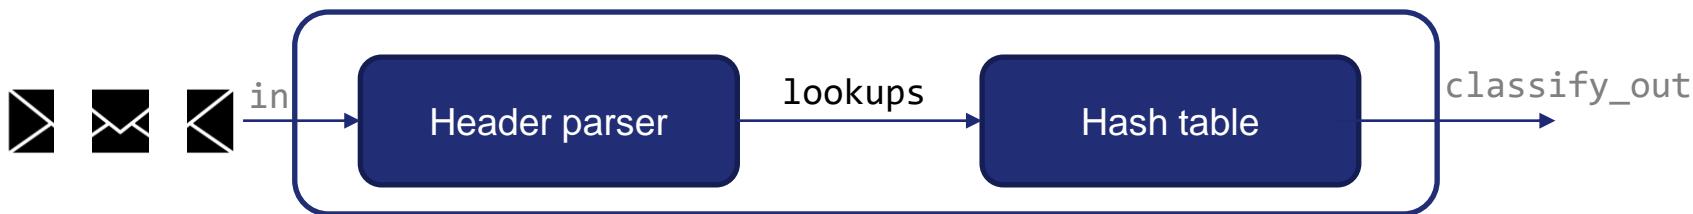
Compiler directive for the data-flow optimization

```
    static hls::stream<hash_tag> lookups;
```

static used to describe the module's internal state

```
    parser(in, lookups);
    hash_table(lookups, classify_out /* ... */);
}
```

Function invocation = hardware module instantiation



# Legacy HLS: simple parser state-machine

```
void parser(axi_data_stream& in, hls::stream<hash_tag>& out)
{
#pragma HLS pipeline
    static enum { IDLE, FIRST, REST } state = IDLE;
    static hash_tag ret;

    axi_data flit;
    ...
}
```

Accept the next flit while processing previous ones

Where we are within the packet

Output may be built over multiple invocations

# Legacy HLS: simple parser state-machine

```
void parser(axi_data_stream& in, hls::stream<hash_tag>& out) {  
...  
    switch (state) {  
        case IDLE:  
            if (in.empty() || out.full())  
                return;  
  
            in.read_nb(flit);  
            ret = extract_headers_flit1(flit);  
            state = flit.last ? IDLE : FIRST;  
            if (flit.last)  
                out.write_nb(ret);  
            break;  
    ...  
}
```

The diagram illustrates the flow of control and data processing in the parser state-machine. It uses callout boxes and arrows to map specific code snippets to their functional components:

- An arrow points from the `if (in.empty() || out.full()) return;` statement to a box labeled "Check for available inputs, output space".
- An arrow points from the `in.read_nb(flit);` statement to a box labeled "Update internal state from input".
- An arrow points from the `out.write_nb(ret);` statement to a box labeled "Output".
- A long arrow originates from the `case IDLE:` label and points to a box labeled "First flit of the packet".

# Legacy HLS: issues

- Static variables make it difficult to reuse code
- Limited use of classes
- Repetitive code handling streams



*There are three great virtues  
of a programmer: **Laziness,**  
**Impatience** and **Hubris**.*

Larry Wall

Creator of the Perl programming language

---

# ntl HLS methodology

# How to build reusable data-flow element pattern?

- Basic elements
  - C++ objects for each data-flow element
  - State kept as member variables
  - `step()` method implements functionality
  - Inline methods embedded in the caller
  - All interfaces are `hls::stream` (members/parameters)
- Reuse with customization via function objects / lambdas
- Composed through aggregation: reusable sub-graph.

# Networking Template Library (ntl)

Class library of packet processing building blocks.

Category	Classes
Header processing elements	pop/push_header, push_suffix
Data-structures	array, hash_table
Scheduler	scheduler
Basic elements	map, scan, fold, dup, zip, link
Specialized stream wrappers	pack_stream, pfifo, stream<Tag>
Control-plane	gateway

# Networking Template Library (ntl)

Class library of packet processing building blocks.

Category	Classes
Header processing elements	<code>pop/push_header</code> , <code>push_suffix</code>
Data-structures	<code>array</code> , <code>hash_table</code>
Scheduler	<code>scheduler</code>
Basic elements	<code>map</code> , <b><code>scan</code></b> , <b><code>fold</code></b> , <code>dup</code> , <code>zip</code> , <code>link</code>
Specialized stream wrappers	<code>pack_stream</code> , <b><code>pfifo</code></b> , <code>stream&lt;Tag&gt;</code>
Control-plane	<code>gateway</code>

# Example: scan and fold

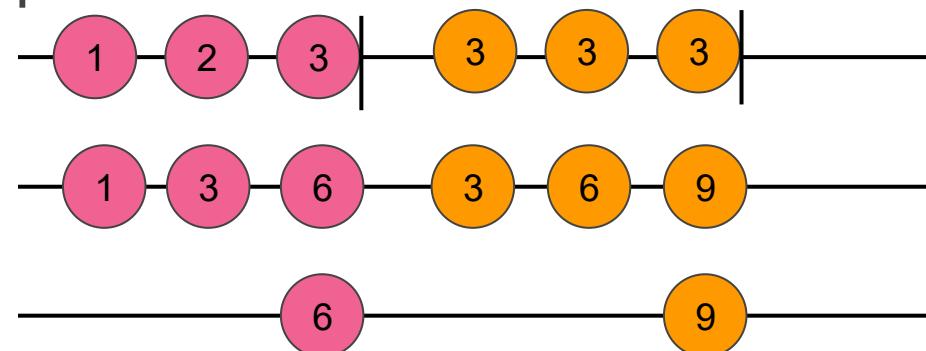
Common operators in functional and reactive programming.

Modified to reset state for every packet.

Input stream:

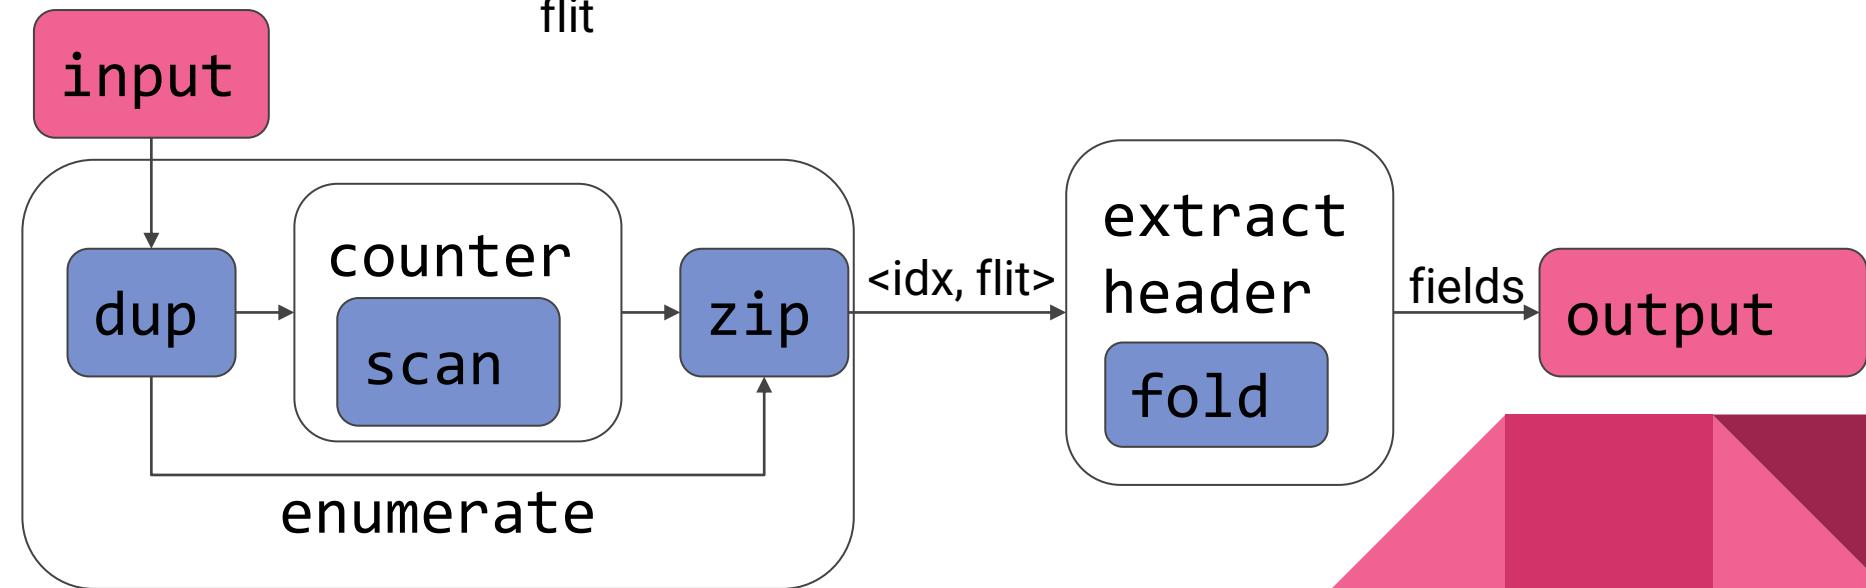
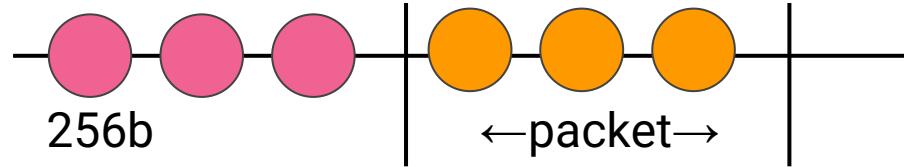
```
scan.step(input, plus())
```

```
fold.step(input, plus())
```



Reduce boilerplate hls::stream handling

# Fold & scan usage: parser example



# Parser class with ntl

```
class parser {  
public:  
    void step(axi_data_stream& in);  
    hls::stream<hash_tag> out;  
  
private:  
    ntl::enumerate<ntl::axi_data> _enum;  
    extract_metadata _extract;  
};
```

Compiler synthesizes the module from the step function

We can instantiate output FIFO within the class to simplify its use

Instantiate sub-modules

# Parser step function

```
void parser::step(axi_data_stream& in)
{
#pragma HLS dataflow
    _enum.step(in);
    _extract.step(_enum.out);
    ntl::link(_extract.out, out);
}
```

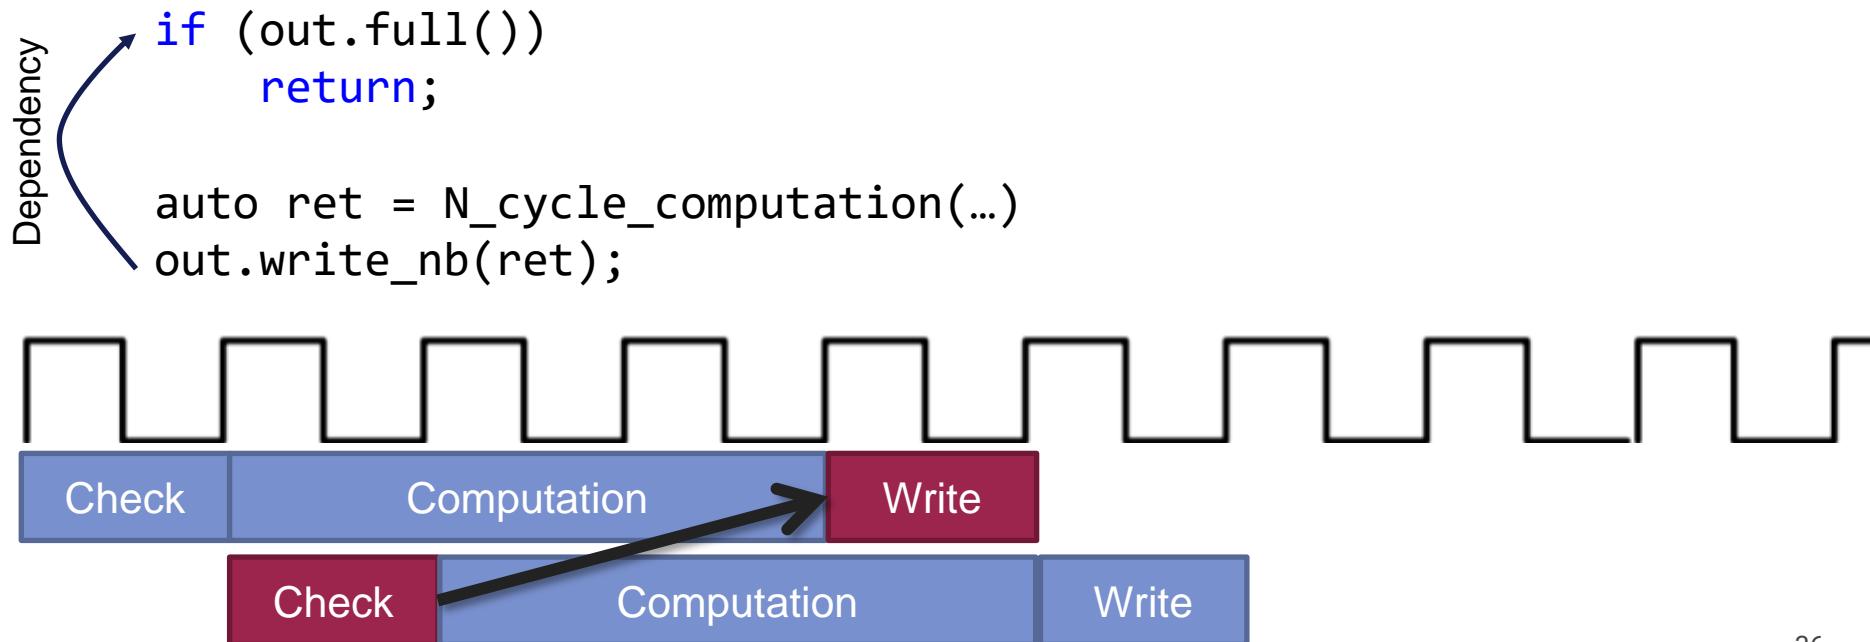
# Top function wrapper

```
void firewall_top(axi_data_stream& in, bool_stream& classify_out)
{
    #pragma HLS dataflow
    static firewall f;
    f.step(in, classify_out);
}
```

static only used once in the  
top function wrapper

# Pipeline dependencies

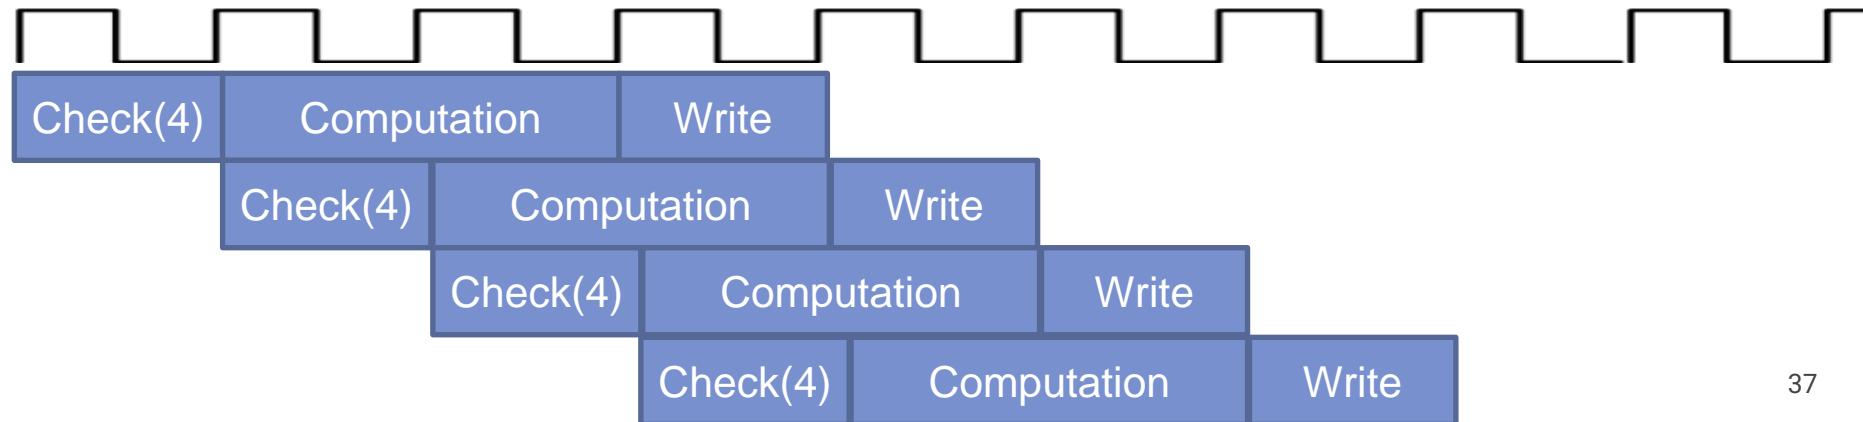
```
#pragma HLS pipeline
```



# Pipeline dependencies

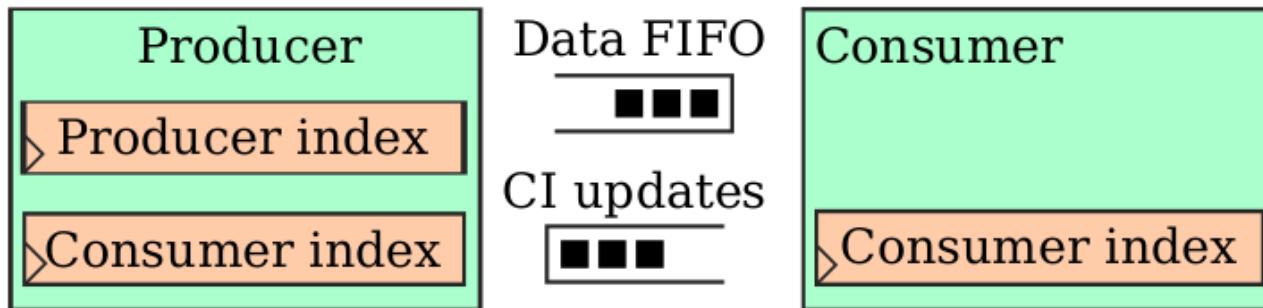
```
#pragma HLS pipeline
```

```
if (out.has_room_for(N+1))
    return;
auto ret = N_cycle_computation(...)
out.write_nb(ret);
```



# Programmable-threshold FIFO

`hls::stream replacement`



# Programmable-threshold FIFO

```
template <typename T, size_t stream_depth>
class programmable_fifo {
public:
    bool write_nb(const T& t);
    bool read_nb(T& t);
    bool full();
    bool empty();
private:
    /* producer and consumer state */
    ...
    hls::stream<T> _stream;
}
```

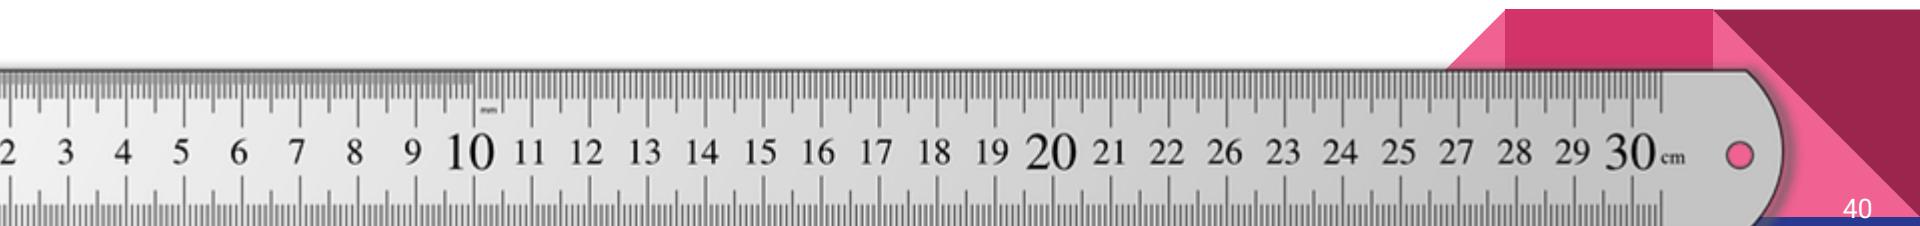
Use compiler directive to  
inline into callers

# Evaluation

- How does nt1 compare against legacy HLS, P4?
- Can we build a relatively complex application with nt1?

Targeting Mellanox Innova Flex SmartNIC

- Xilinx Kintex UltraScale XCKU060 FPGA
- 216.25 MHz clock rate



# Stateless UDP firewall example

Use on-chip hash-table to classify packets.

	Thpt.	Latency	LUTs	FFs	BRAM	LoC
HLS/nt1	72 Mpps	25 cycles	5296	7179	12	218
HLS legacy	72 Mpps	16 cycles	4087	4287	12	593
P4 (SDNet 2018.2)	108 Mpps	211 cycles	34531	49042	193	92

# Stateless UDP firewall example

Use hash-table to classify packets.

All exceed line rate (59.5 Mpps)

	Thpt.	Latency	LUTs	FFs	BRAM	LoC
HLS/nt1	72 Mpps	25 cycles	5296	7179	12	218
HLS legacy	72 Mpps	16 cycles	4087	4287	12	593
P4 (SDNet 2018.2)	108 Mpps	211 cycles	34531	49042	193	92

# Stateless UDP firewall example

Use hash-table to classify packets.

x2.7 less lines of code compared to legacy

	Thpt.	Latency	LUTs	FFs	BRAM	LoC
HLS/nt1	72 Mpps	25 cycles	5296	7179	12	218
HLS legacy	72 Mpps	16 cycles	4087	4287	12	593
P4 (SDNet 2018.2)	108 Mpps	211 cycles	34531	49042	193	92

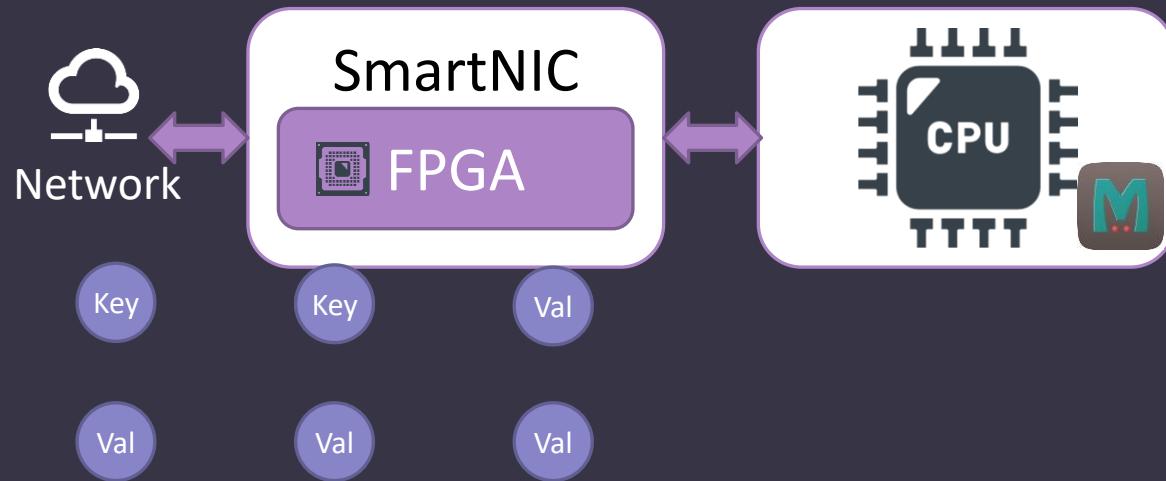
# Stateless UDP firewall example

Use hash-table to classify packets.

	Thpt.	Latency	LUTs	FFs	BRAM	LoC
HLS/nt1	72 Mpps	25 cycles	5296	7179	12	218
HLS legacy	72 Mpps	16 cycles	4087	4287	12	593
P4 (SDNet 2018.2)	108 Mpps	211 cycles	34531	49042	193	92

nt1 requires more LoC, but improves latency & area

# A Key-Value Store Cache



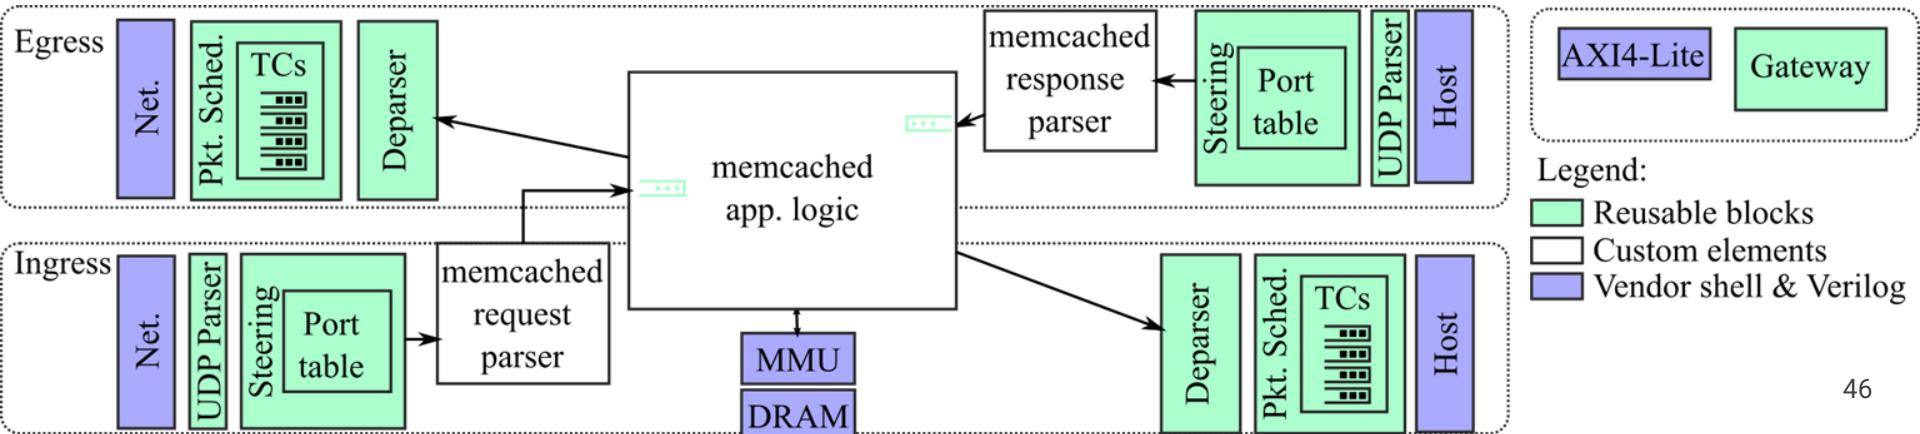
Eran, H., Zeno, L., Tork, M., Malka, G., & Silberstein, M. [NICA: An infrastructure for inline acceleration of network applications](#). In 2019 USENIX Annual Technical Conference (USENIX ATC 19).

# Key-value store cache

Processes 16-byte GET hits at 40.3 Mtps.

For 75% hit rate: 9x compared to CPU-only.

Uses: hash tables, header processing, scheduler, control plane, programmable FIFOs, ...



# More information & related work

For more information

- Our paper: [Design Patterns for Code Reuse in HLS Packet Processing Pipelines](#), FCCM '19.
- [Productive parallel programming for FPGA with HLS](#), Johannes de Fine Licht and Torsten Hoefer, ETH.
- Xilinx's [Introduction to FPGA Design with Vivado High-Level Synthesis](#) and [Vivado Design Suite User Guide – High-Level Synthesis](#)

Related work

- [https://github.com/Xilinx/HLS\\_packet\\_processing](https://github.com/Xilinx/HLS_packet_processing)  
Xilinx (unofficial) packet processing library builds parsers with boost::mpl.
- [Module-per-Object: a human-driven methodology for C++-based high-level synthesis design](#), Silva et al., FCCM '19.

# Conclusion

What does it take to write packet processing hardware in C++ HLS?

Try out ntl: <https://github.com/acsl-technion/ntl> 

Thank you!

Questions?