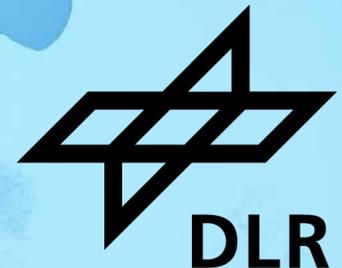
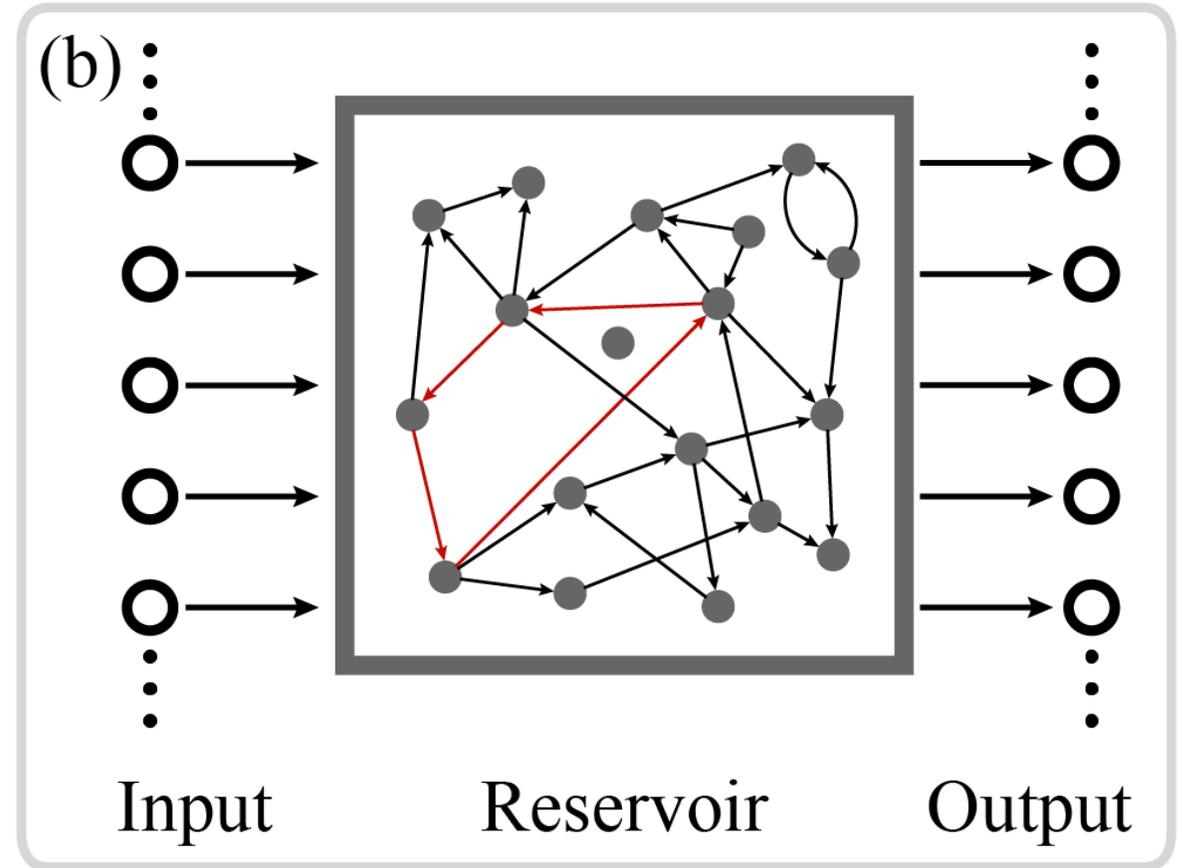
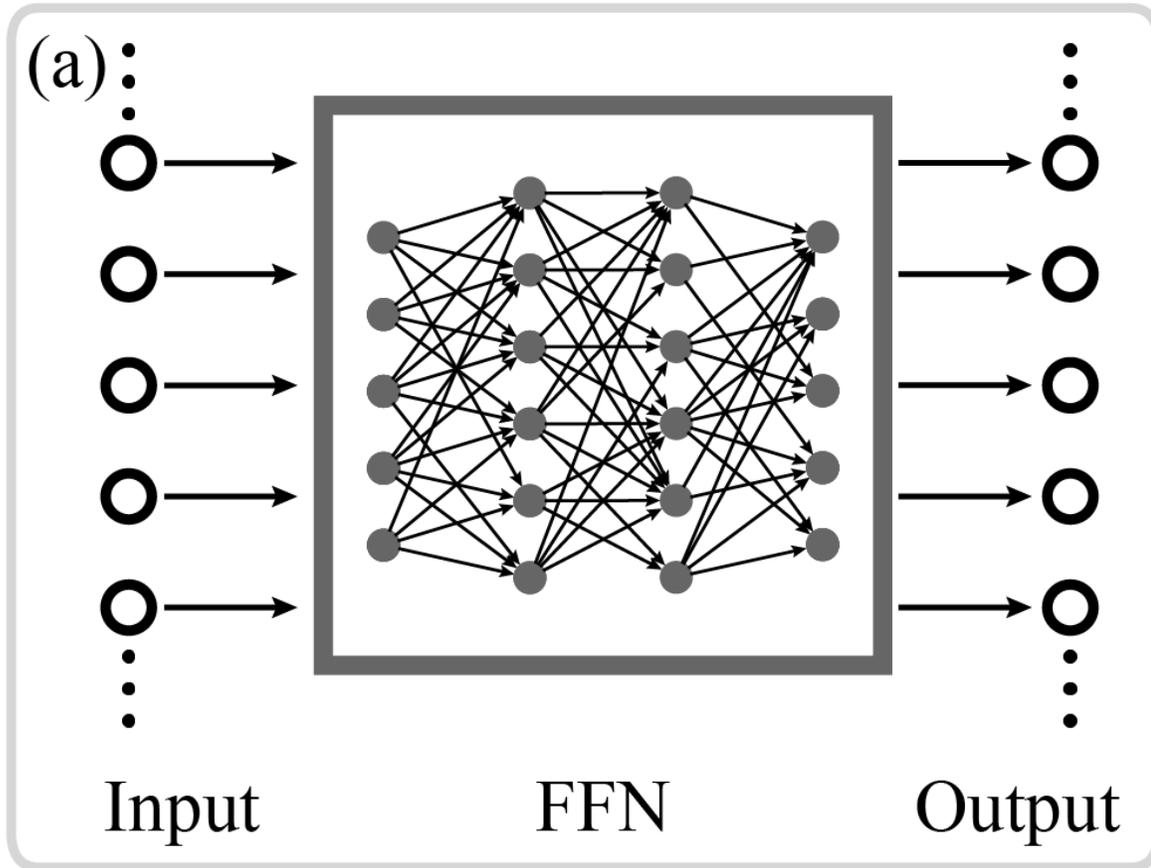


FPGA IMPLEMENTATION OF NGRC FOR PREDICTING DYNAMICAL SYSTEMS

By João Folhadela and Daniel Köglmayr – DLR KI-EXE
2nd FPGA Developers' Forum (FDF) meeting (CERN)



Introduction – Reservoir Computing (RC)



Reservoir Computing

Optimize W_{out} via Ridge regression. Minimize loss:

$$\sum_{t=-T}^0 \| W_{out} \mathbf{r}(t) - \mathbf{y}_R(t) \|^2 + \beta \| W_{out} \|^2$$

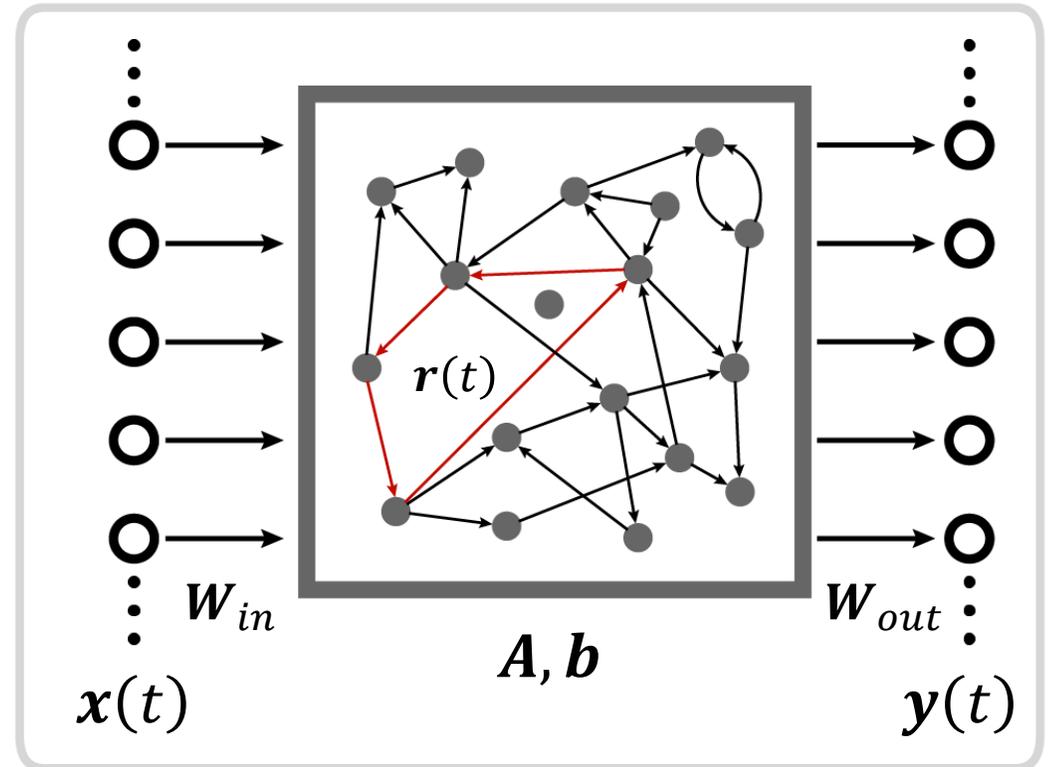
$\mathbf{y}_R(t)$: known real output

β : regularization parameter

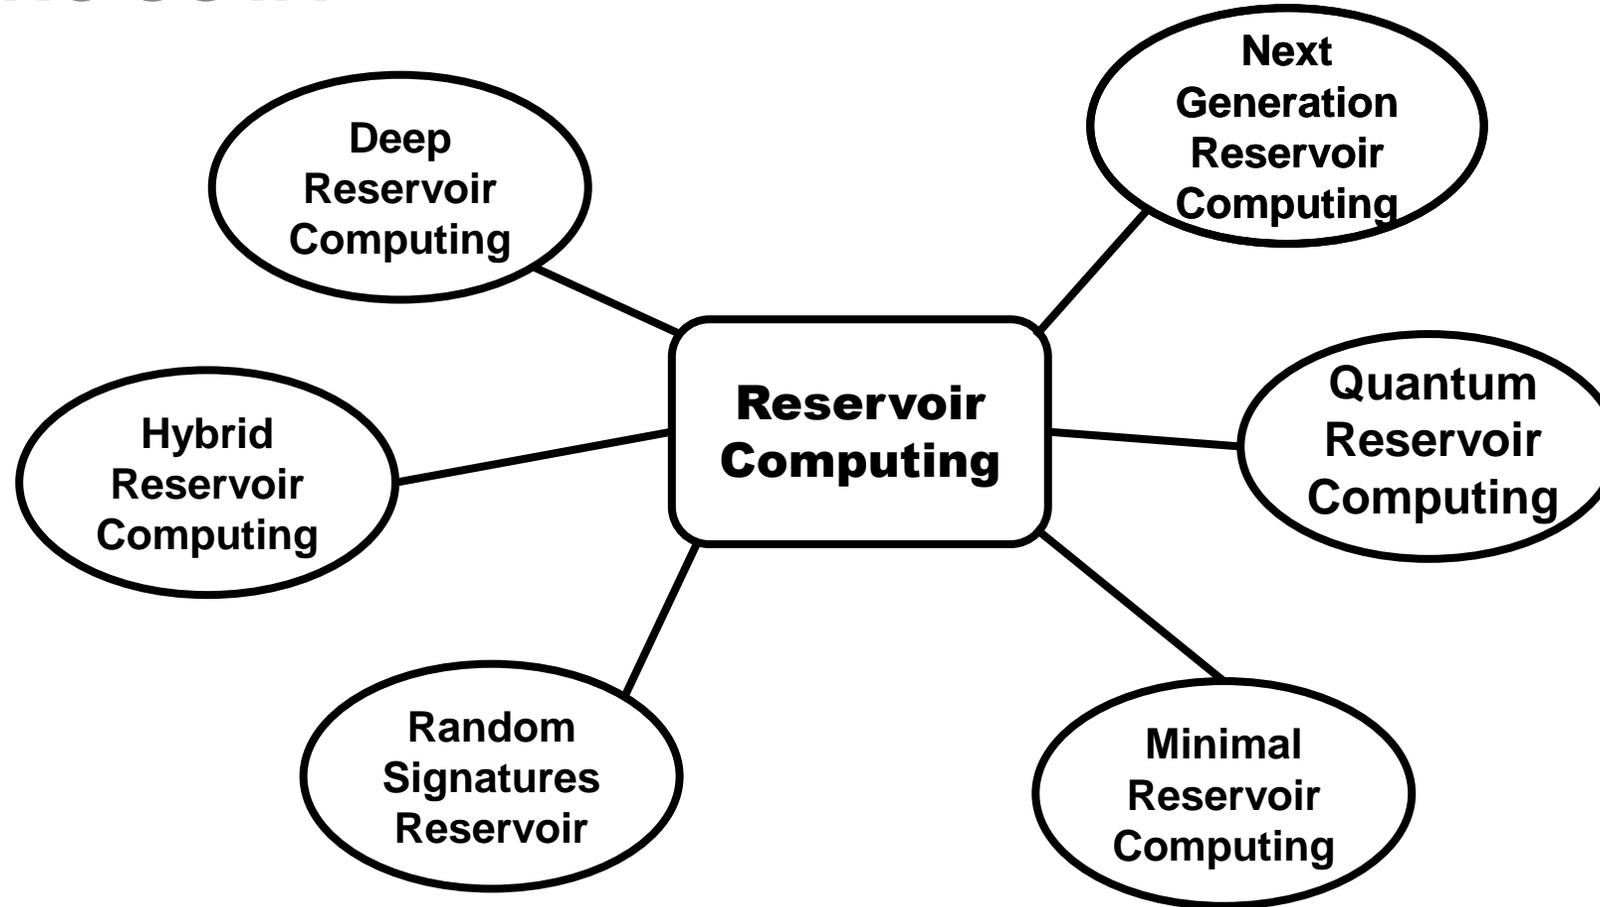
W_{out} is then obtained via:

$$W_{out} = (\mathbf{r}^T \mathbf{r} + \beta \mathbf{I})^{-1} \mathbf{r}^T \mathbf{y}$$

An analytic solution is usually much faster to compute than gradient descent.

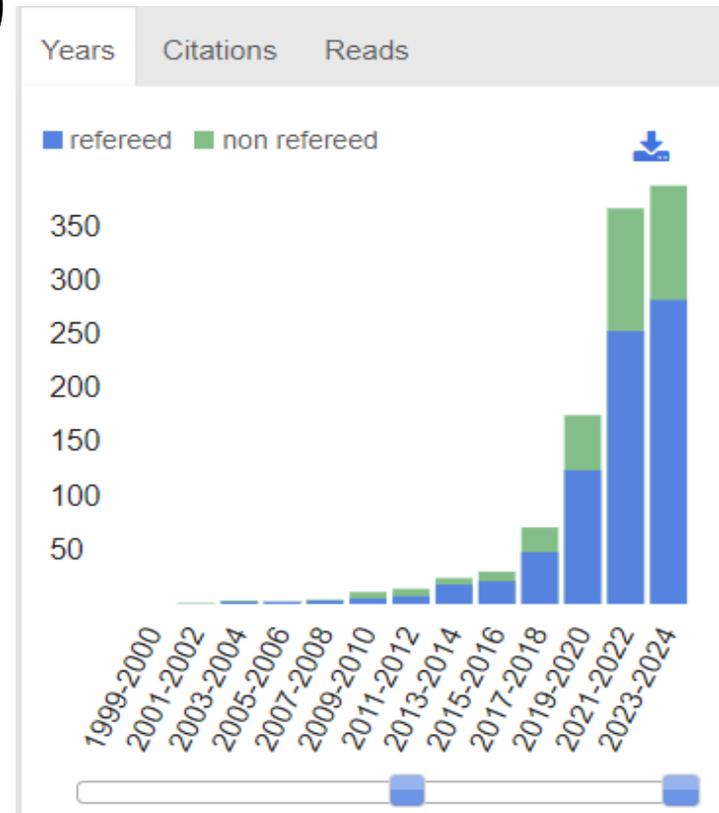


RC SOTA

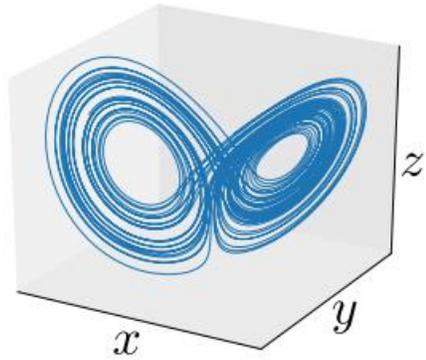


All methods have two basic architectural components:

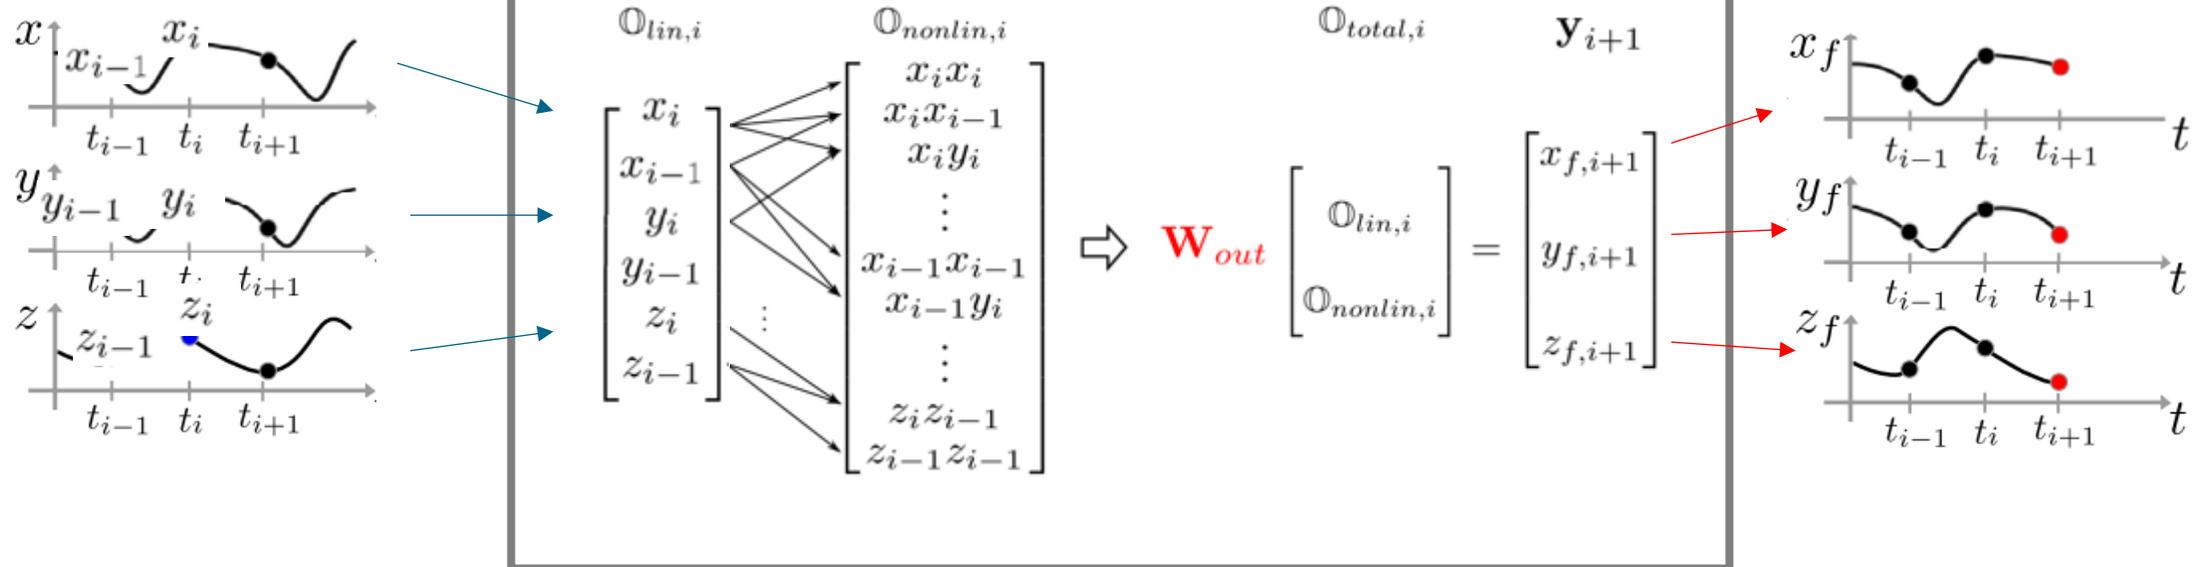
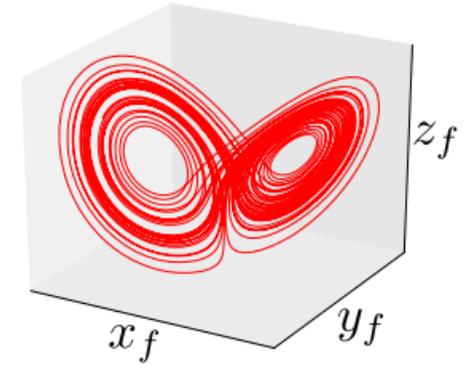
- Nonlinear dimensionality expansion of the input data
- Regression based mapping on the target data



Introduction – Next Generation Reservoir Computing (NGRC)



- Efficient
- Less training data required



Introduction – NGRC on Lorenz Attractor



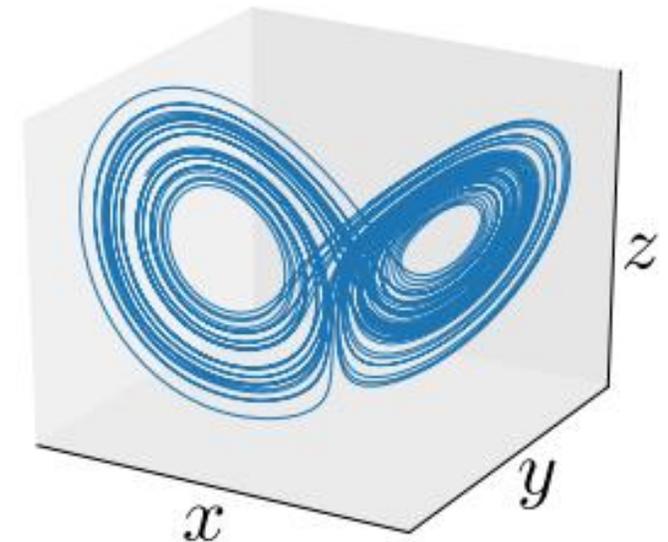
- We train with Python's numpy's linalg solver
- Some parameters:

$$W_{out} = (r^T r + \beta I)^{-1} r^T Y_{train}$$

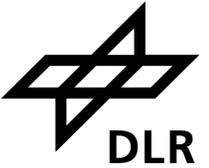
Train Steps	k	β	W_{out} shape
400	2	9×10^{-2}	27x3

k – timesteps β - regularization parameter

- Then inference is performed on FPGA
 - Transfer W_{out} and first linear state to FPGA memory
 - Prediction loop



Implementation – The NGRC prediction loop



1. Creating the Reservoir

- Read linear state from memory – current and previous timestep
- Non-linear expansion – **multiplying and powering the linear state**

2. Predicting

- Matrix multiplication, several vector dot products – **Multiply and Accumulate** (MACs)

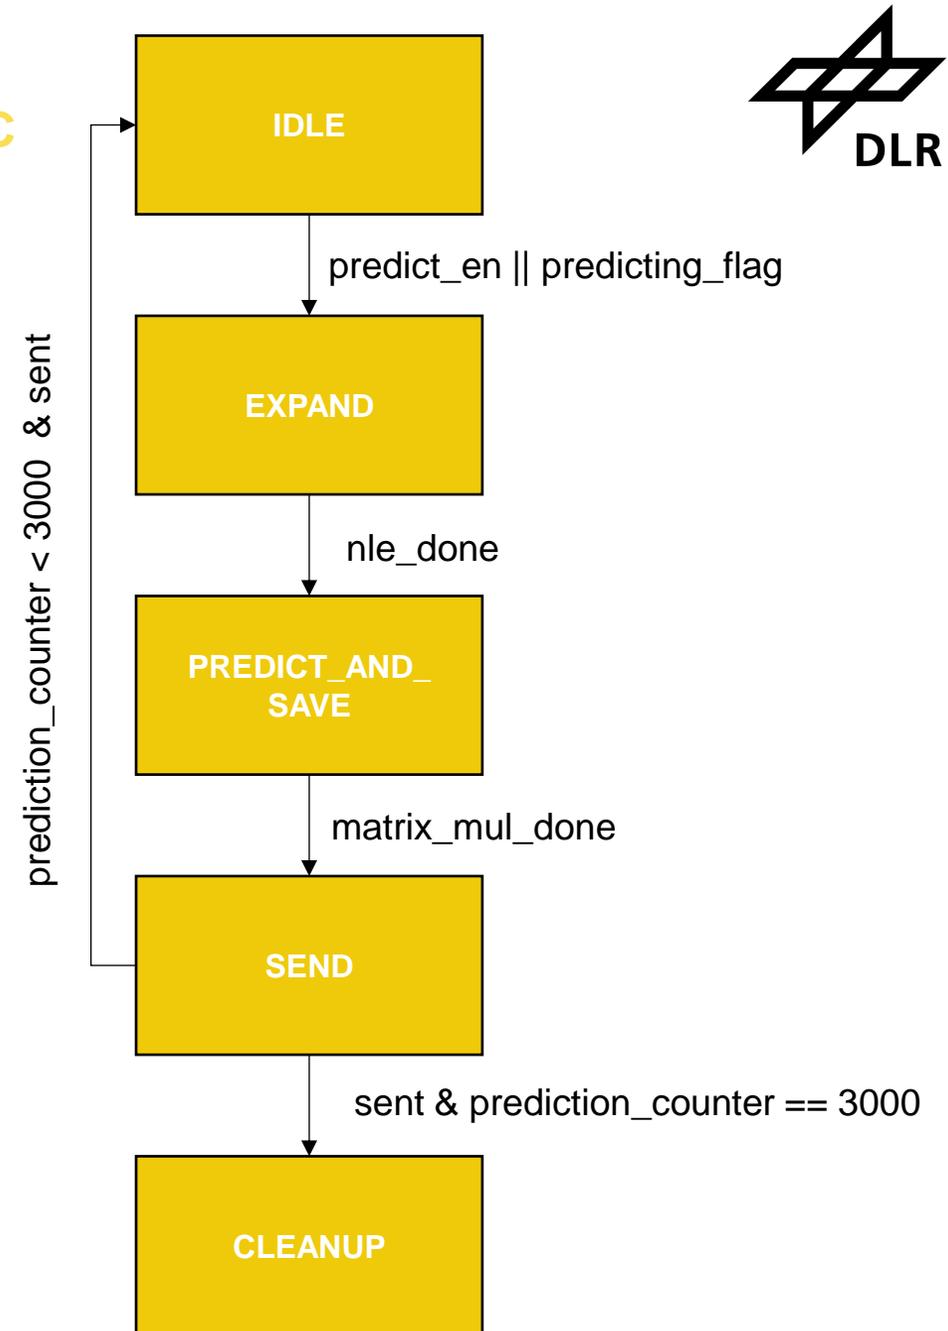
3. Save prediction

In red – arithmetics are performed in float 32 with Vivado floating-point IP

Implementation – The FSM

- **IDLE**
 - Sets linear state
 - Waits for prediction loop to begin/continue
- **EXPAND**
 - Performs the non-linear expansion
- **PREDICT_AND_SAVE**
 - Calculates prediction
 - Saves it
- **SEND**
 - Send the current prediction via UART
- **CLEANUP**
 - Resets all counters and variables

NGRC



Implementation – The FSM

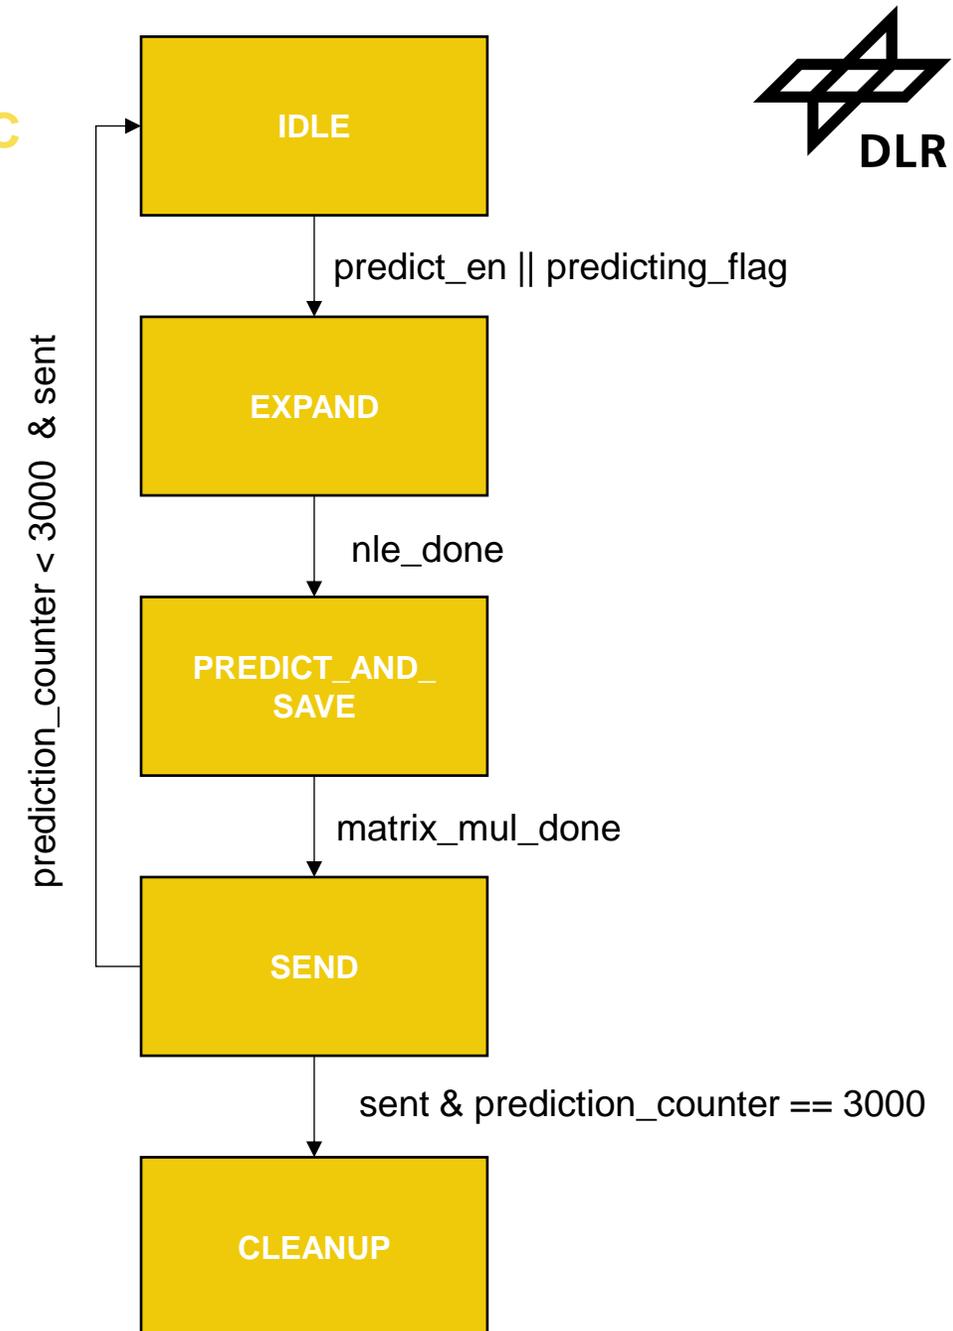


- Implemented as a module that expands a vector of size 6 into size 27 (6+21)
- Each multiplication is computed in parallel, with 32-bit float numbers:
 - Vivado Floating-point IP
 - *Multiply* operation
 - *Full DSP usage* (21 multiplications x 2 DSP = 42 DSPs)
 - Maximum Latency used (8 clock cycles)
- Concatenate the input vector with each of the multiplication results.

Implementation – The FSM

- **IDLE**
 - Sets linear state
 - Waits for prediction loop to begin/continue
- **EXPAND**
 - Performs the non-linear expansion
- **PREDICT_AND_SAVE**
 - Calculates prediction
 - Saves it
- **SEND**
 - Send the current prediction via UART
- **CLEANUP**
 - Resets all counters and variables

NGRC



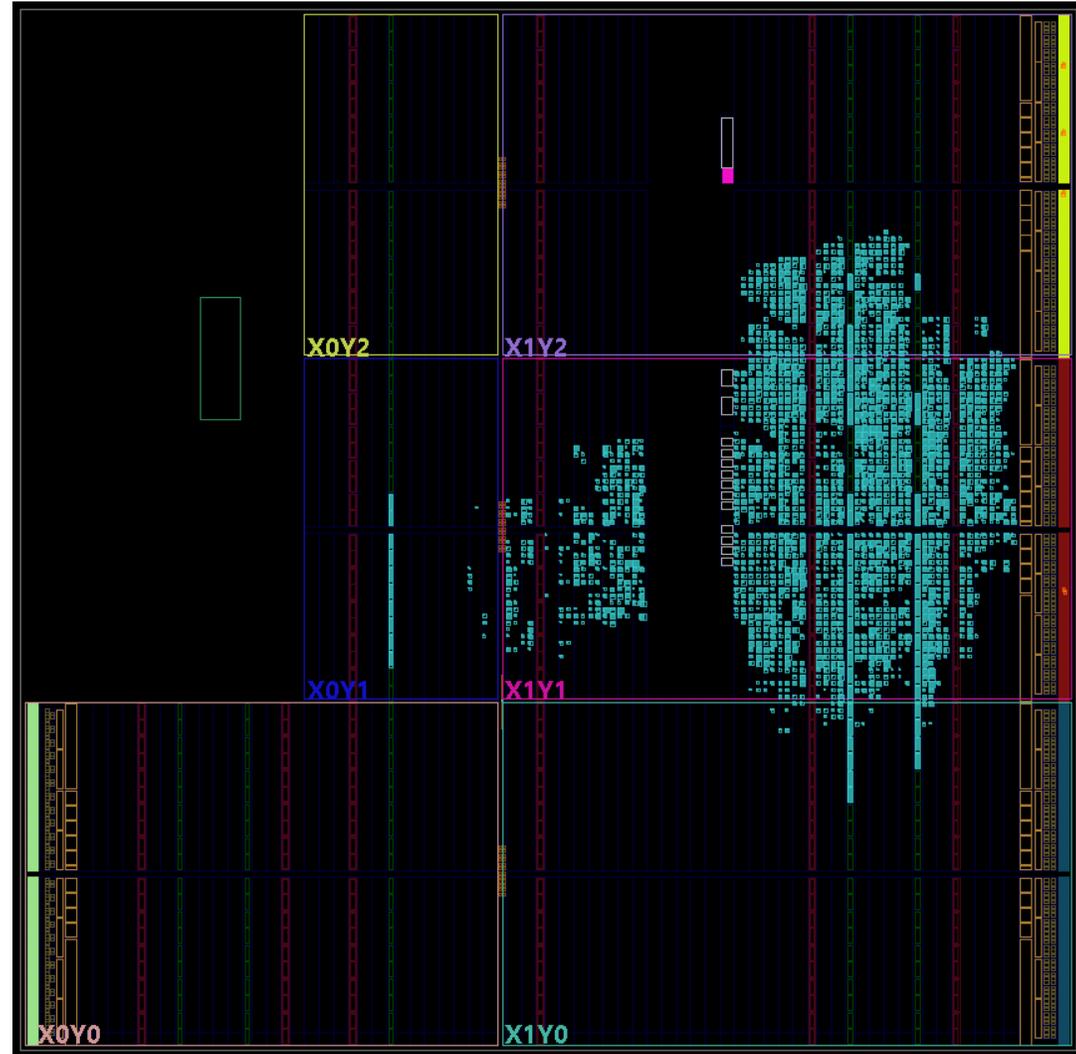
Implementation – The FSM



- Perform matrix multiplication of the reservoir state (a vector) with the weights, W_{out} :
 - 3 MACs implemented in parallel with Vivado Floating-point IP that compute:
 - *Fused Multiply-Add* operation
 - *Full DSP usage* (3 MACs x 4 DSP = 12 DSPs)
 - 10 clock cycles latency used
 - 10 latency x 27 accumulations = 270 clock cycles in total
- Save prediction (to update next linear state)

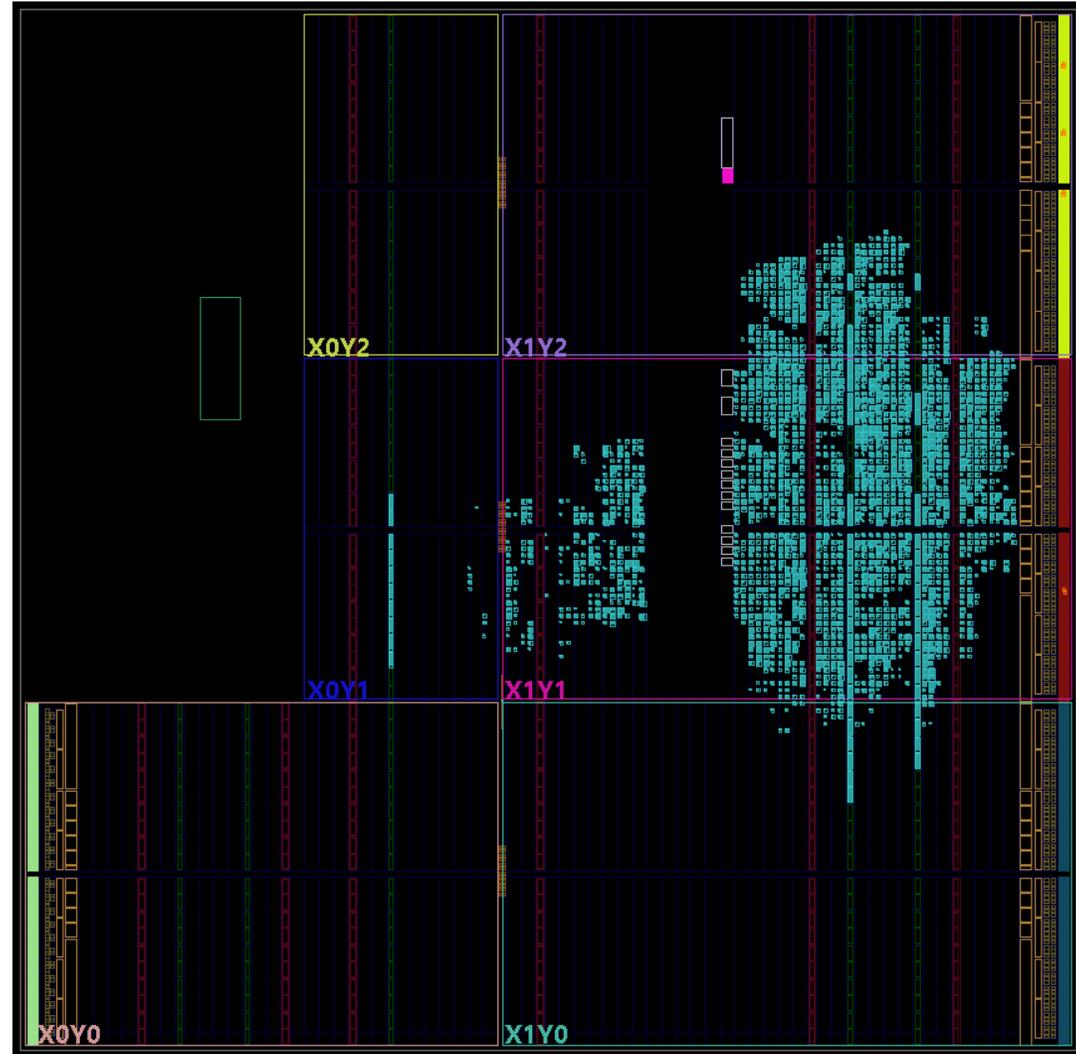
Resource usage

Resource	Used / Available
LUT	5,132 / 53,200
LUTRAM	270 / 17,400
FF	6,475 / 106,400
DSP	54 / 220



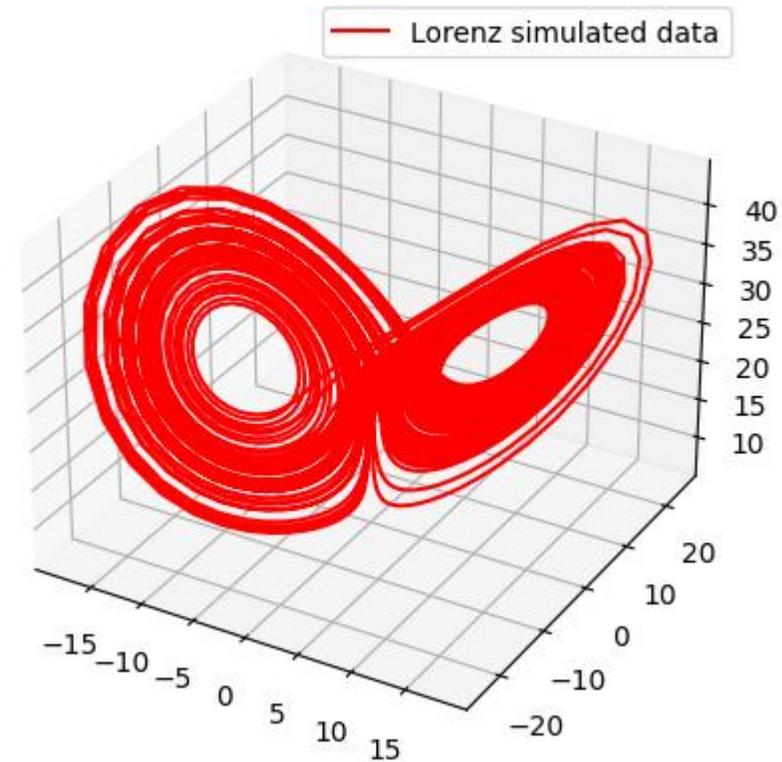
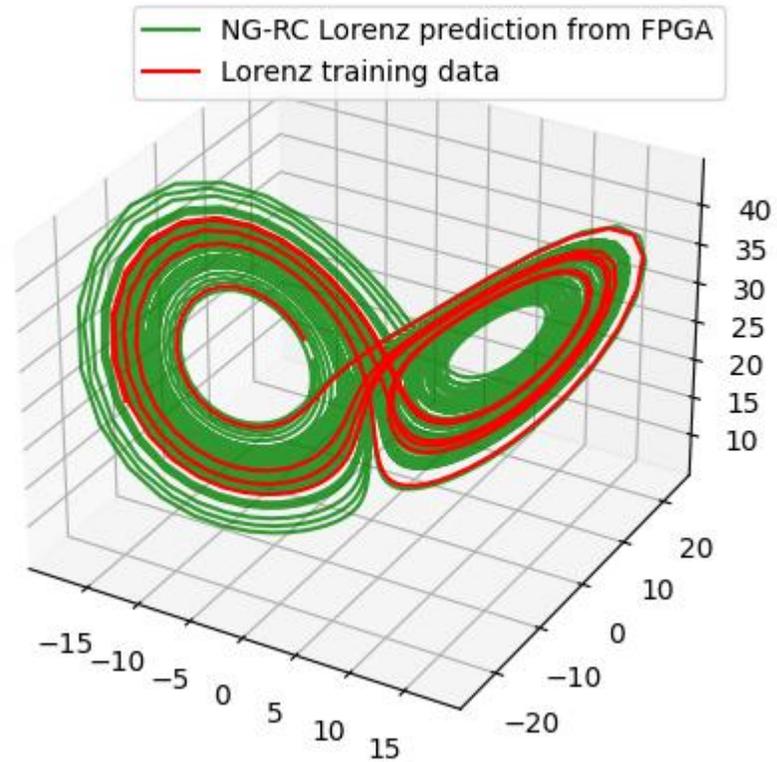
Other reports

Property	Value
On-Chip Power	0.294 W
Dynamic	0.186 W
Static	0.109 W
Clock Speed	125 MHz
Worst Negative Slack	1.3 ns



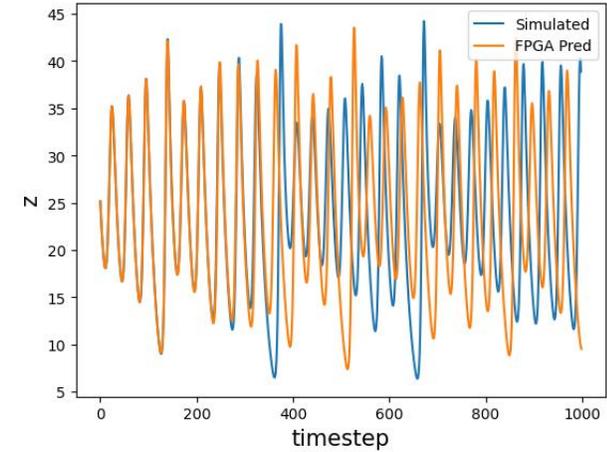
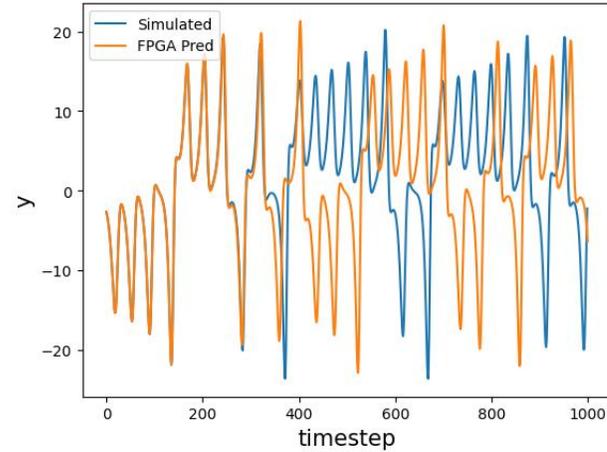
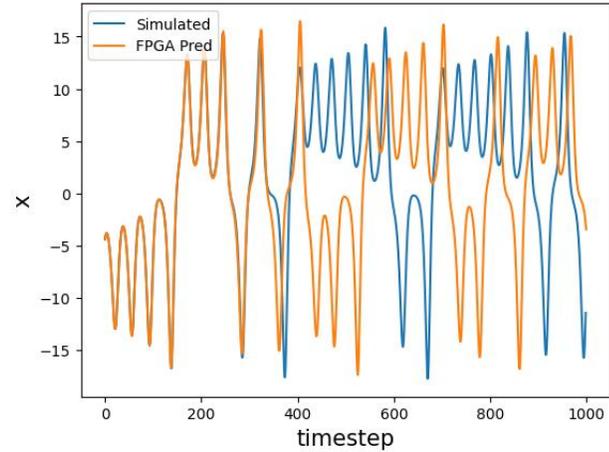
Evaluation – Simulated vs. Predicted Lorenz

Prediction over 3000 timesteps

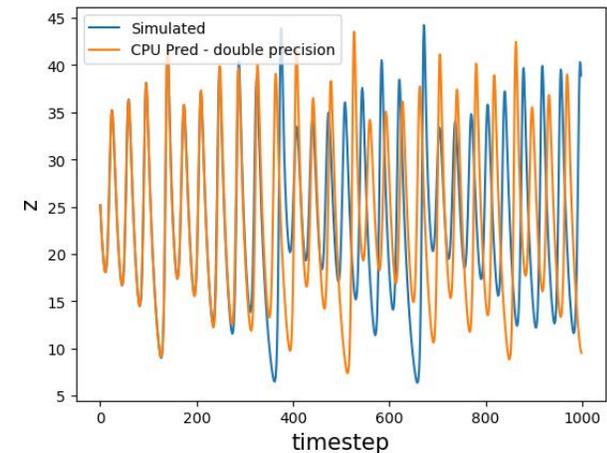
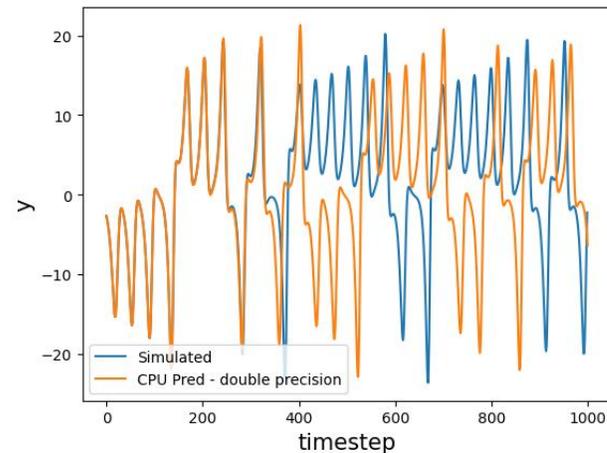
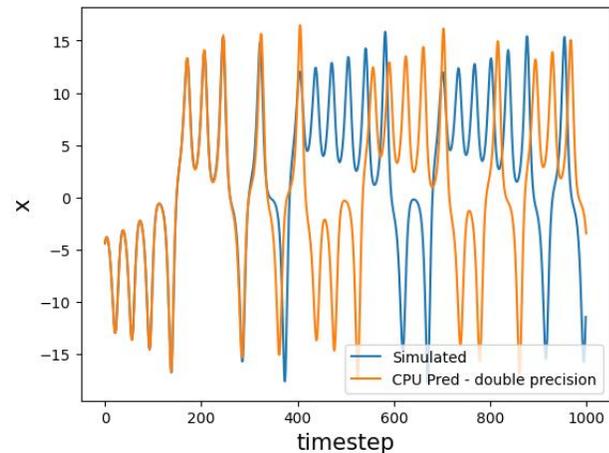


Evaluation – Simulated vs. Predicted Lorenz

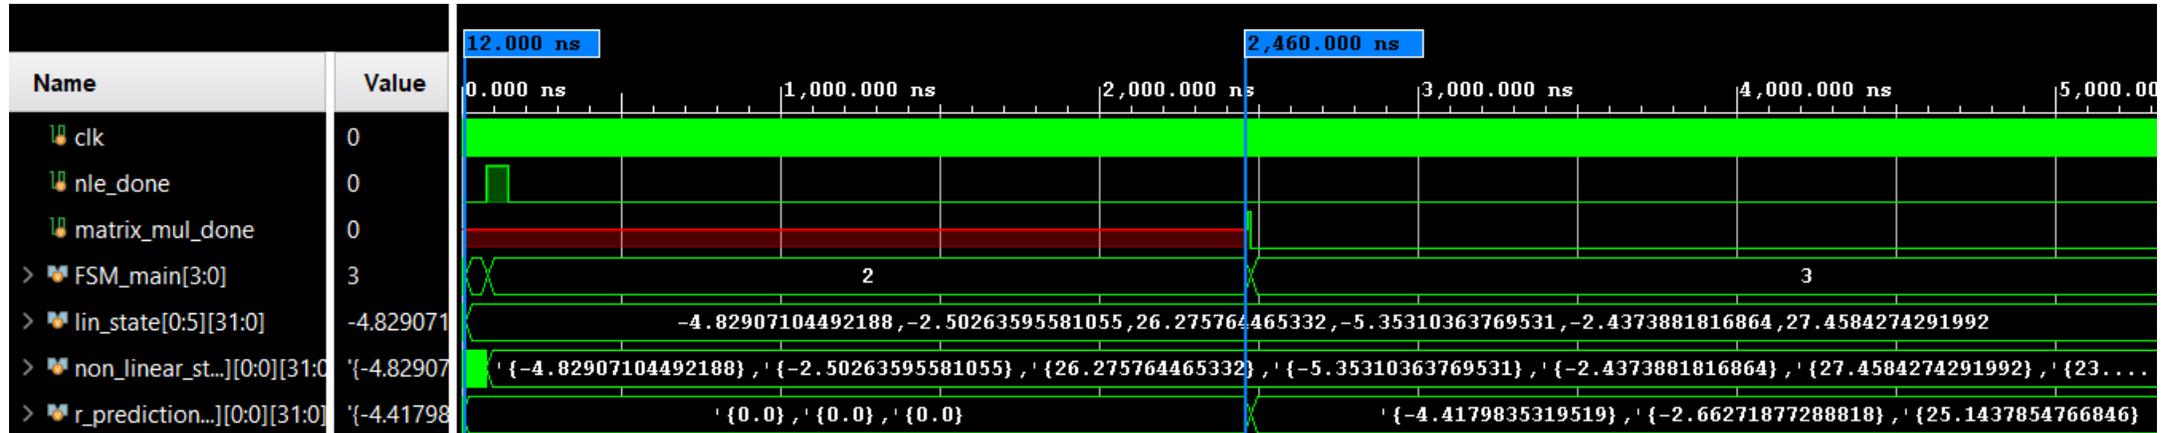
On FPGA



On CPU – with double precision



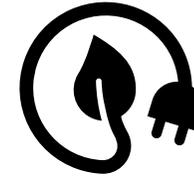
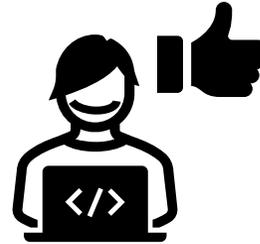
Latency Evaluation – Simulated vs. Predicted Lorenz



	FPGA	CPU
Device	Zynq7000-20	i7-1365U 13 th Gen
Clock Speed	125 MHz	4.8 GHz
Inference time	2.5 μ s	50 μ s

Evaluation – Predicting the Lorenz attractor

- Easy to implement

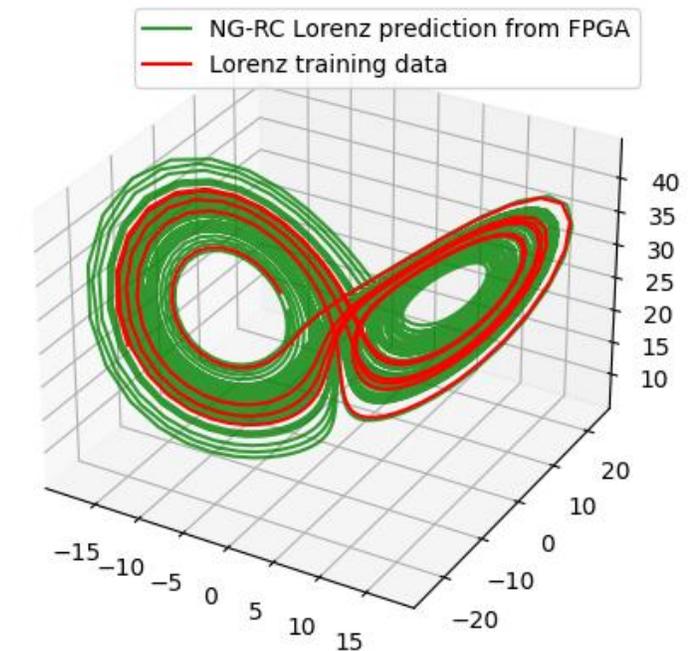


- Low power

- Fast



And real applications?

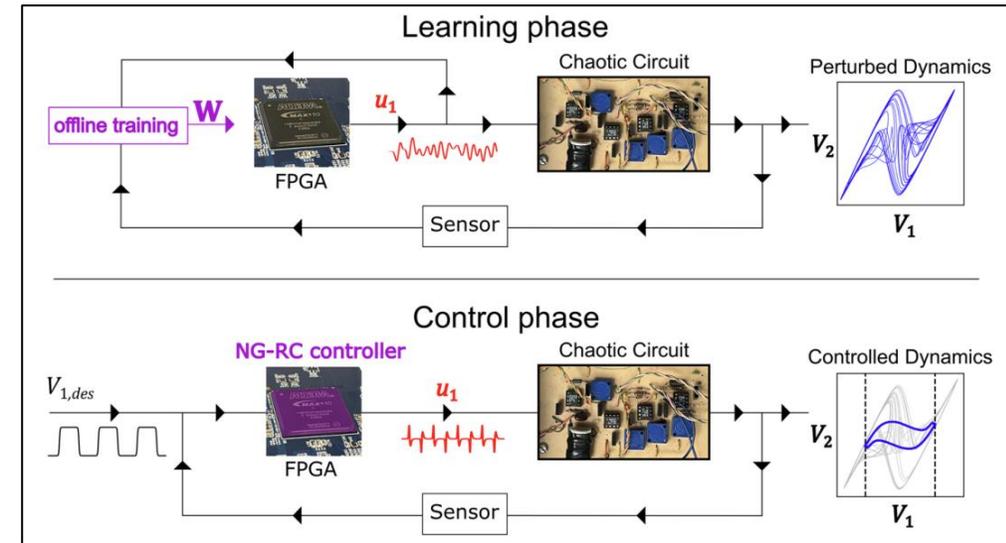


Applications of NGRC

■ Control of chaotic circuit dynamics

Kent, R.M. et. al *Controlling chaos using edge computing hardware*

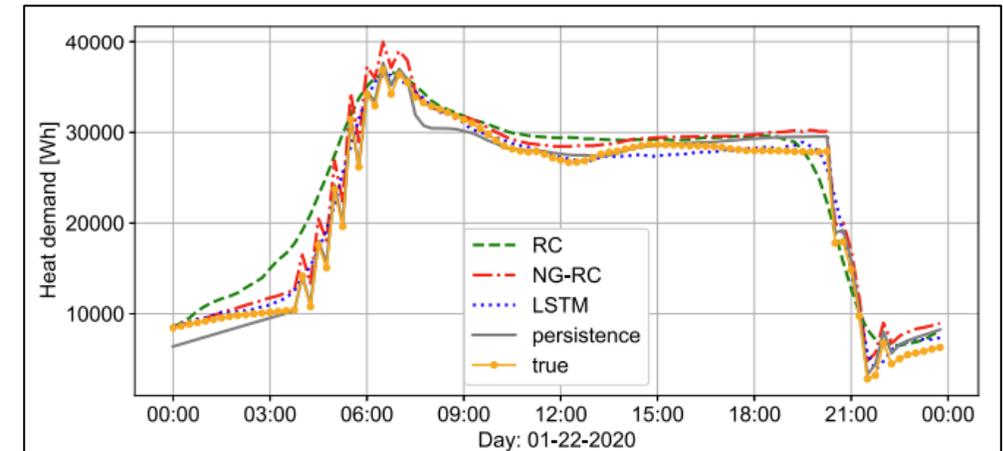
Nat Commun 15, 3886 (2024)



■ Energy demand forecasting

Karoline Brucke et. al *Benchmarking reservoir computing for residential energy demand forecasting*

Energy and Buildings 314 (2024)



■ Financial markets, weather, costal line erosion, robotic control, etc.

THANK YOU!

Questions?

