

CM1K speed performance

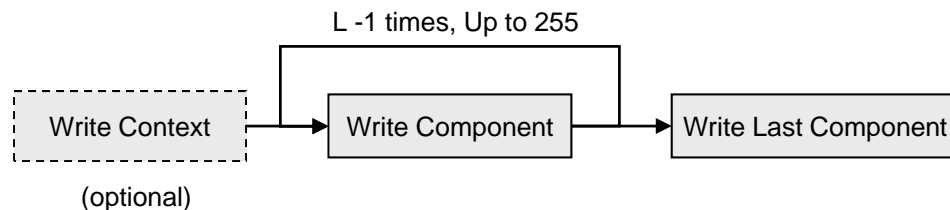


- All 1024 neurons operate in parallel
- Learn and recognize a 256-bytes vector < 11 μ sec @27Mhz
- Save or restore of the 1024 neurons < 10 ms @27Mhz
- Disruptive performance for standard GPU !

Learn and Recognize Ops

< 11 μ s / vector

1) Broadcast a vector of L bytes...




Clock cycles @27Mhz
L=256

(1)+L+3

9.55 μ s


2) ...to learn



19

0.66 μ s

Or 3)...to recognize and response of M first neurons

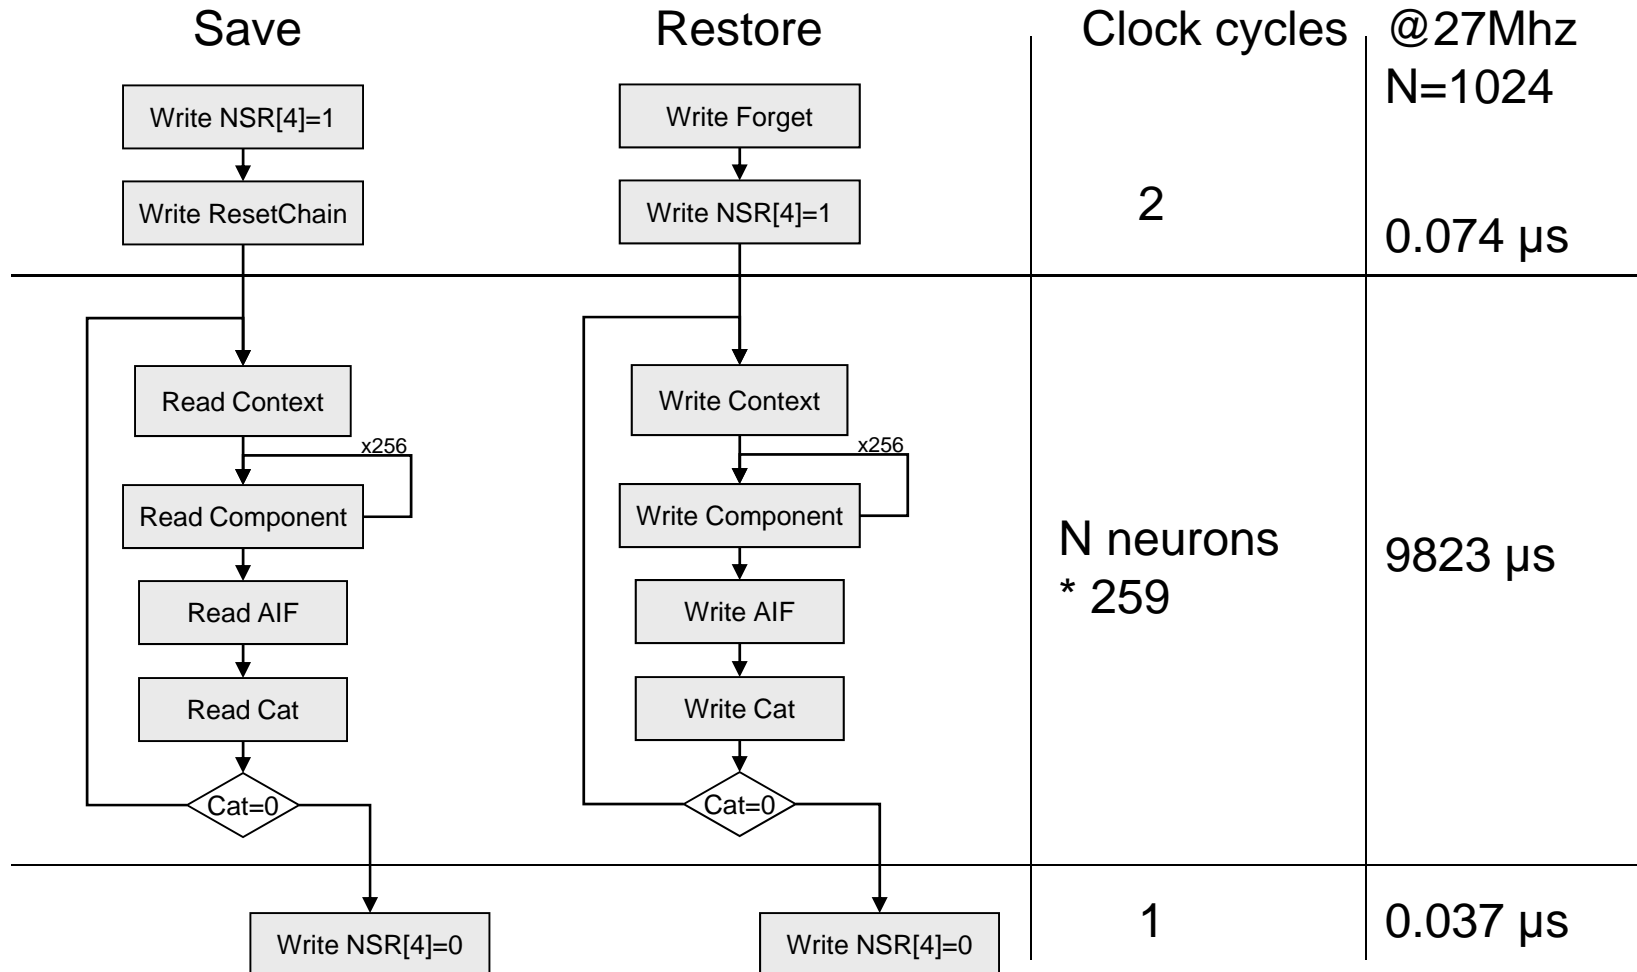


M * (18+19)

M=1

1.33 μ s

Save and Restore ops (9.8 ms/chip @27Mhz)

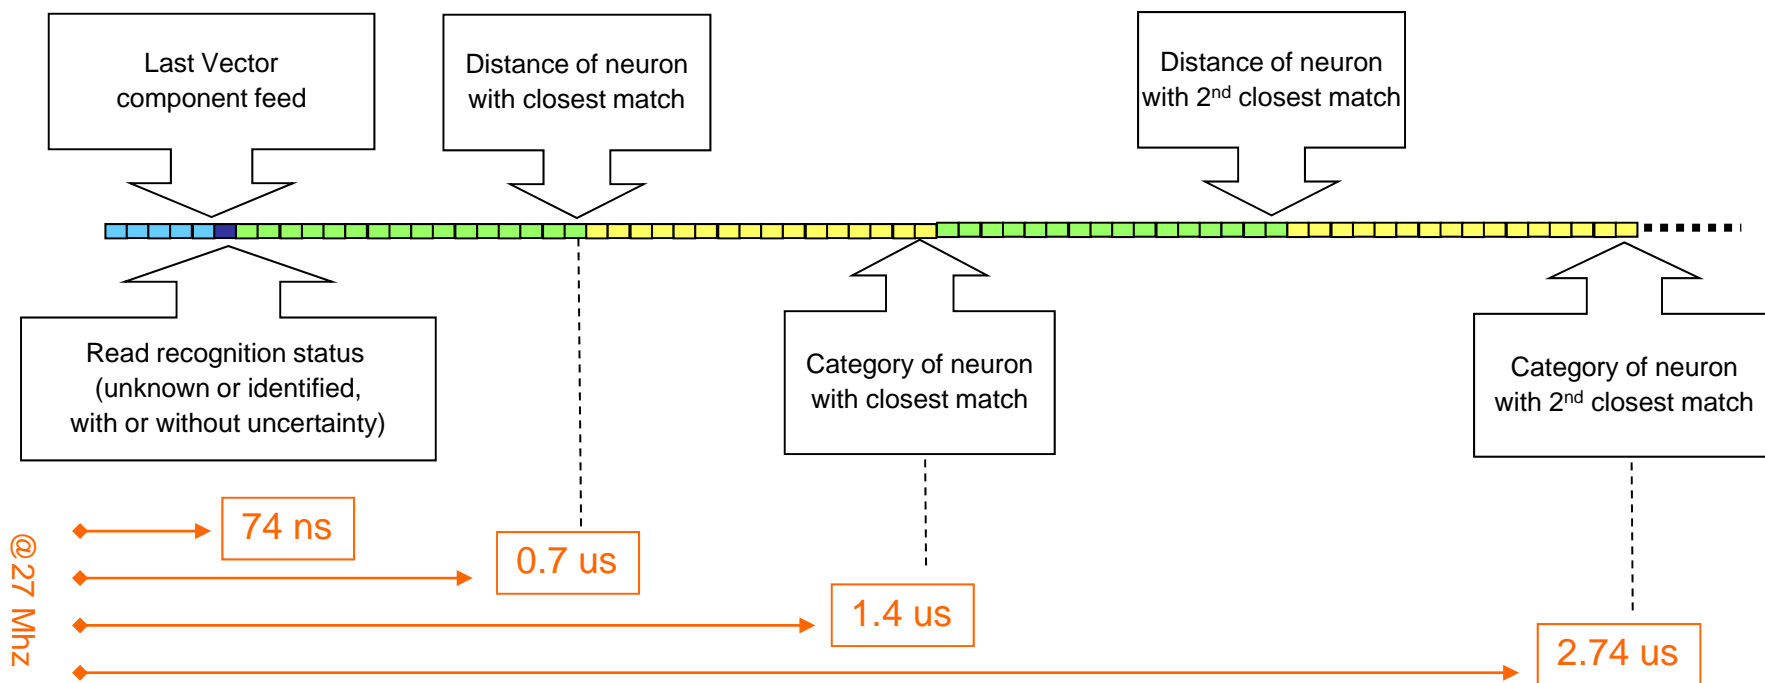


CM1K Register Access Time

Addr	Name	Description	Default	LR mode		SR mode	
0x00	NCR	[6:0]=Neuron Context [7]= Norm [15:8]=Neuron identifier[23:16]	0x0001			W (1)	R (1)
0x01	COMP	Component	0x00	W (1)		W (1)	R (1)
0x02	LCOMP	Last Component	0x00	W (1 if no neuron is committed, 3 otherwise)			
0x03	IINDEX	Component Index	0x00	W (1)		W (1)	
0x03	DIST	Distance	0xFFFF		R (18)		R (1)
0x04	CAT	Category[14:0] Degenerated flag[15]	0xFFFF	W (1 if ID, 19 otherwise)	R (3 if ID, 19 otherwise)	W (1)	R (1)
0x05	AIF	Active Influence Field	0x4000			W (1)	R (1)
0x06	MINIF	Minimum Influence Field	0x0002	W (1)	R (1)		R (1)
0x07	MAXIF	Maximum Influence Field	0x4000	W (1)	R (1)		
0x0A	NID	Neuron Identifier[16:0]	0x0000		R (1)		R (1)
0x0B	GCR	[7]= Norm [6:0]= Global Context	0x0001	W (1)	R (1)		
0x0C	RESET CHAIN	Point to the 1 st neuron in the chain				W (1)	
0x0D	NSR	Network Status Register	0x00	W (1)	R (1)	W(1)	
0x0F	FORGET	Forget (clear the knowledge and some global registers)*		W (1)			
0x0F	NCOUNT	LR mode: committed neurons SR mode: neuron index			R (1)		R(1)

3 levels of recognition

- From the feed of the last vector component...
 - Level 1: Recognition Status = +1 cc (ex: anomaly detection)
 - Level 2: Closest distance = +18 cc (ex: tracking)
 - Level 3: Best match category = +19 cc (ex: identification)
- Readout of subsequent distance and category = + 37 cc



@27Mhz

Recognition/seconds

RBF

Vector length (bytes)	Vector reco* (clk cycles)	Time** (usec)	Reco/sec
256	293	10.84	92,250
128	165	6.10	163,934
64	101	3.73	268,096

*Reco = read nearest distance and associated category

KNN

Vector length (bytes)	Vector 5NN* (clk cycles)	Time** (usec)	Reco/sec
256	347	12.83	77,942
128	219	8.01	124,843
64	154	5.69	175,746

*KNN (K=5)= read 5 nearest distances

Thanks to CM1K parallelism, No need to know how many neurons are committed to fill this table!

Comparison with DSP for Manhattan distance calculation

Pattern length	256
Number of neurons, N	1,024

L1 distance calculation per neuron	CM1K
Write components (less last one)	255
Write last component + flag if distance < AIF	3
Read Distance	18
Memory access V[i]	
Absolute Subtraction + Branch	
Init + Accumulation	
Branch and Write if Dist[i] < AIF[i]	
Branch and Write if MinDist > Dist[i]	
Total cycles for 1 neuron	276

In the CM1K, ALL the neurons calculate their distance in parallel in the 276 clock cycles time lapse

Smallest L1 distance between 1 vector and N neurons

	CogniMem	DSP SHARK 21160	DSP Tiger-SKARK TS101
clock frequency (Mhz)	27	100	300
Clock cycle (ns)	37.04	10.00	3.33
Number of instructions	276	N*774+7	N*278+34
Total cycles	276	792,583	284,706
Total time (usec)	10.22	7925.83	949.02
Ratio		775	93

CM1K is 93 times faster!