Christof Teuscher

Computation with Structured and Unstructured Networks of Emerging Devices

EPFL Workshop on Logic Synthesis and Emerging Technologies, Sep 28-29, 2017

Portland State University
Department of Electrical and Computer Engineering (ECE)
Department of Computer Science (CS)
Systems Science Graduate Program (SySc)
www.teuscher-lab.com | teuscher@pdx.edu

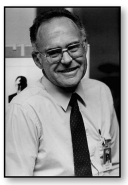
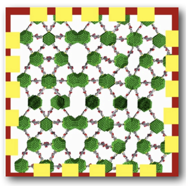
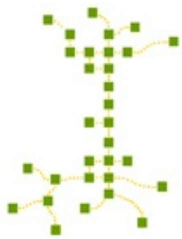

teuscher.:Lab
teuscher-lab.com
Emerging Computing Models and Technologies

Portland State UNIVERSITY

Our Research in a Nutshell

"Our goal is to develop disruptive new computing paradigms and machines that will allow for lasting breakthroughs and open new application domains in the next 5-20 years."

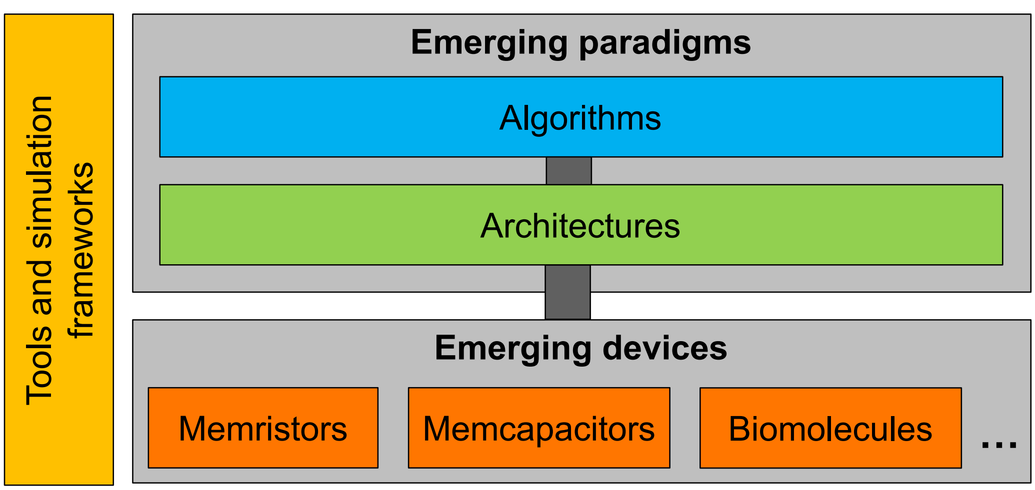

Computing with Structured vs Unstructured Substrates

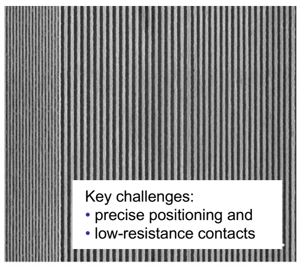

Melosh et al., Science, 2003

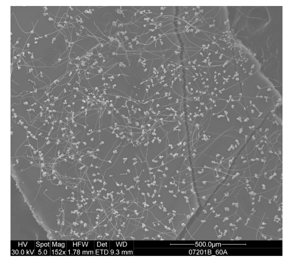
Polyaniline (PANI) conductive polymer, LANL, Wang et al.

Computing with Structured vs Unstructured Substrates

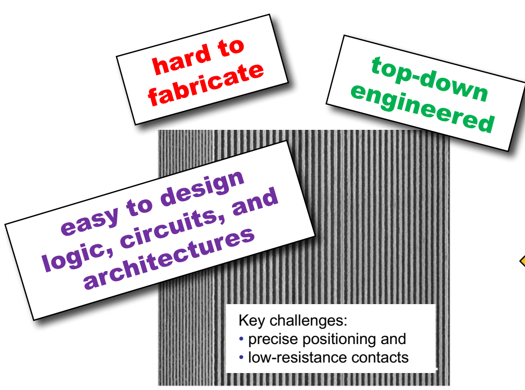

Melosh et al., Science, 2003

Polyaniline (PANI) conductive polymer, LANL, Wang et al.

Christof Teuscher ■ www.teuscher-lab.com ■ Portland State University

Computing with Structured vs Unstructured Substrates

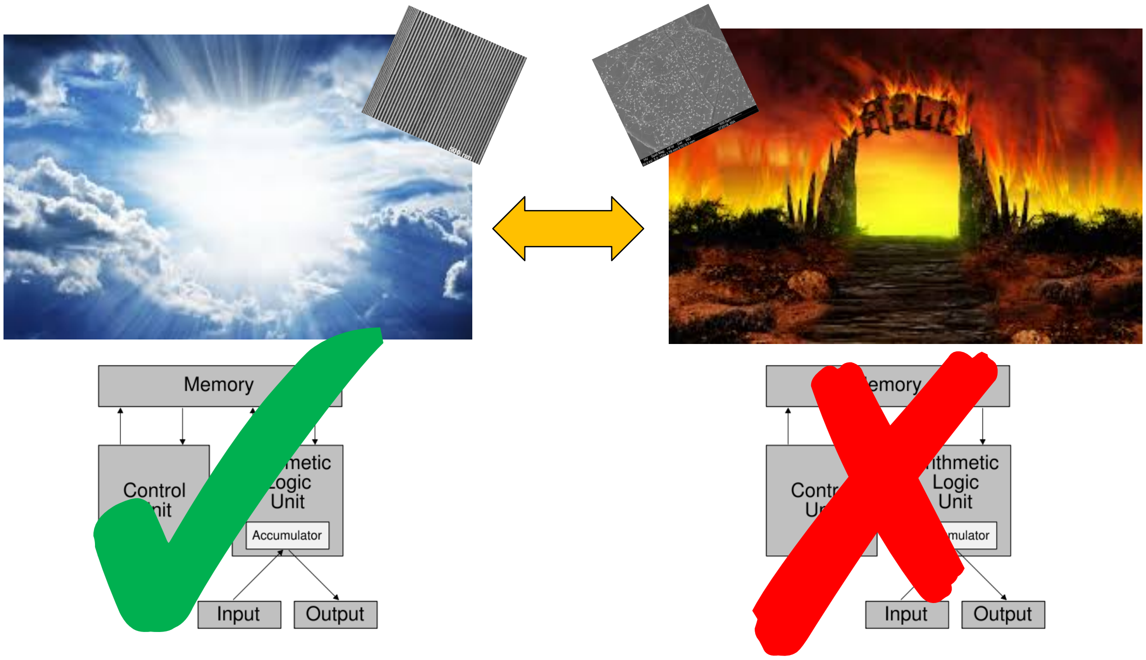

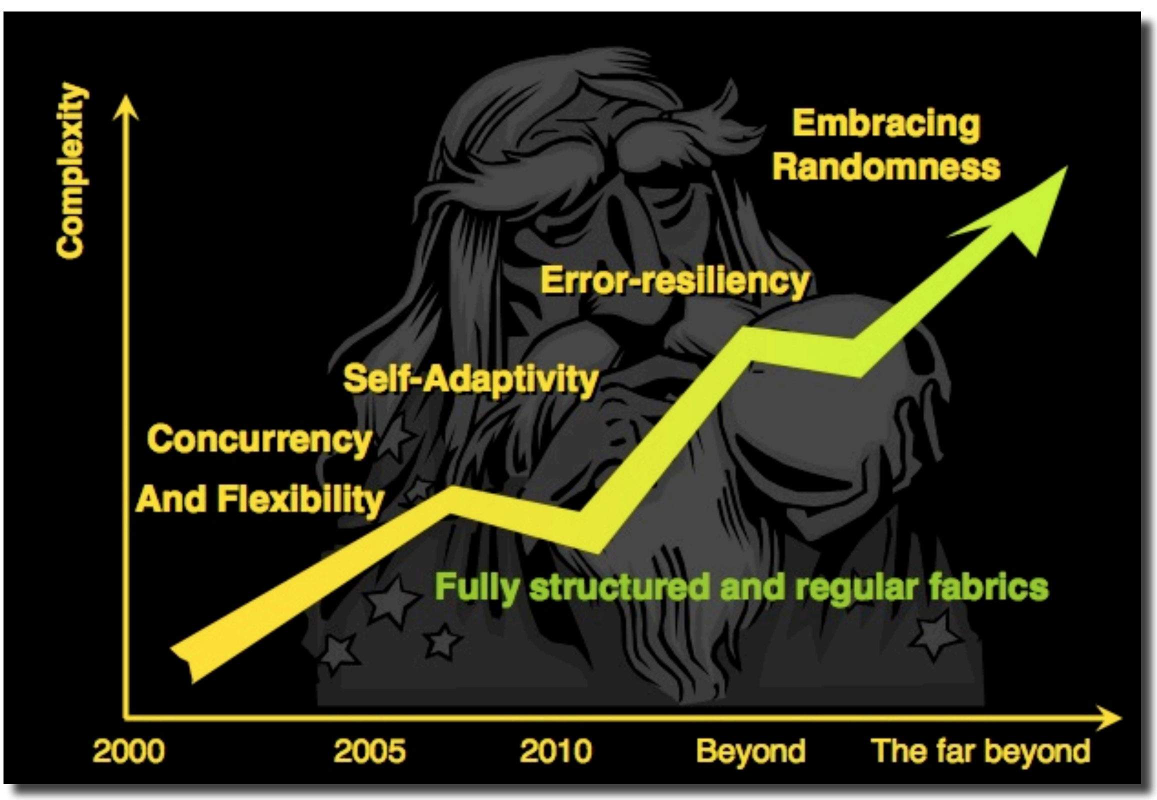

J. Rabey

What's the next big thing in computing? And how do we get there?

As feature-size scaling and "Moore's Law" in CMOS circuits further slow, attention is shifting to computing by non-von Neumann, non-CMOS, and non-Boolean computing models.

IEEE rebooting COMPUTING

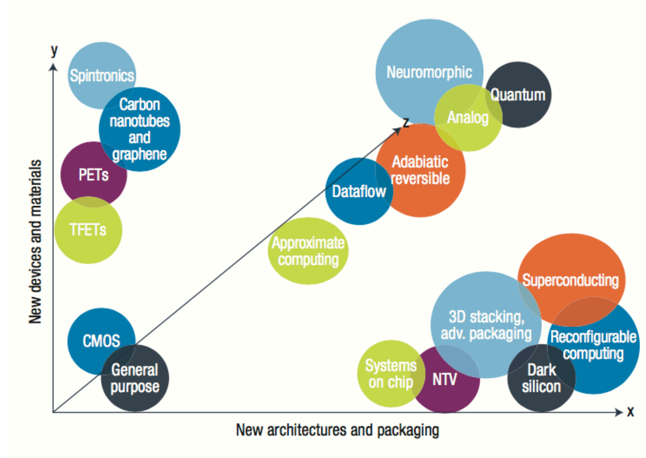

FIGURE 1. Technology scaling options along three dimensions. The graph's origin represents current general-purpose CMOS technology, from which scaling must continue. All the dimensions, which are not mutually exclusive, aim to squeeze out more computing performance. PETs: piezo-electric transistors; TFETs: tunneling field-effect transistors; NTV: near-threshold voltage.

Material Implication

- J. Borghetti, G. S. Snider, P. J. Kuekes, J. J. Yang, D. R. Stewart, and R. S. Williams. **Memristive switches enable stateful logic operations via material implication**. Nature, 464:873–876, 2010.

a

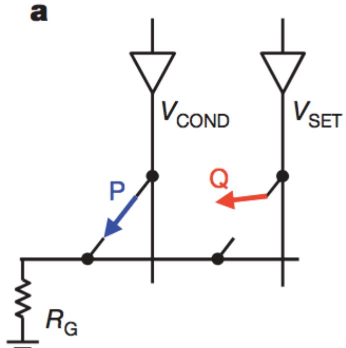

b

$q' \leftarrow p\mathrm{IMP}q$

In p	In q	Out q'
0	0	1
0	1	1
1	0	0
1	1	1

c

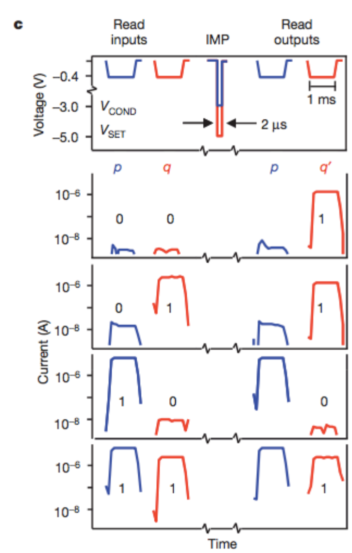

Digital Logic Synthesis for Memristors (1)

- J. Burger, C. Teuscher, and M. Perkowski, **Digital Logic Synthesis for Memristors**, Reed-Mueller Workshop, pp 31-40, 2013.

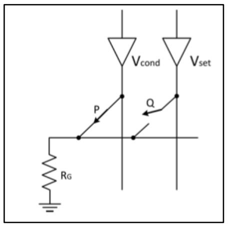

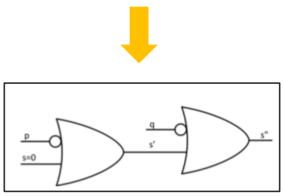

+ **ABC**

Benchmark	Inputs	Pulse Count [10]	Gate count (ABC)	PulseCount (ABC)
exam1_d.pla	3	29	9	12
exam3_d.pla	4	25	10	12
rd53f1.pla	5	30	17	27
rd53f2.pla	5	54	30	57
rd53f3.pla	5	121	18	32
xor5_d.pla	5	121	18	32
con1f1.pla	7	31	11	18
con2f2.pla	7	16	14	19
rd73f1.pla	7	140	104	238
rd73f2.pla	7	253	26	46
rd73f3.pla	7	210	54	104
newill_d.pla	8	6	32	50
newtag_d.pla	8	27	13	21
rd84f1.pla	8	212	109	351
rd84f2.pla	8	330	32	47
rd84f3.pla	8	10	15	23
rd84f4.pla	8	420	134	345
9sym_d.pla	9	420	410	1418
max46_d.pla	9	219	200	427
sao2f1.pla	10	33	56	102
sao2f2.pla	10	31	66	112
sao2f3.pla	10	22	152	380
sao2f4.pla	10	33	125	252
sym10_d.pla	10	1260	346	1172
t481_d.pla	16	320	441	1564

Digital Logic Synthesis for Memristors (2)

- J. Burger, C. Teuscher, and M. Perkowski, **Digital Logic Synthesis for Memristors**, Reed-Mueller Workshop, pp 31-40, 2013.

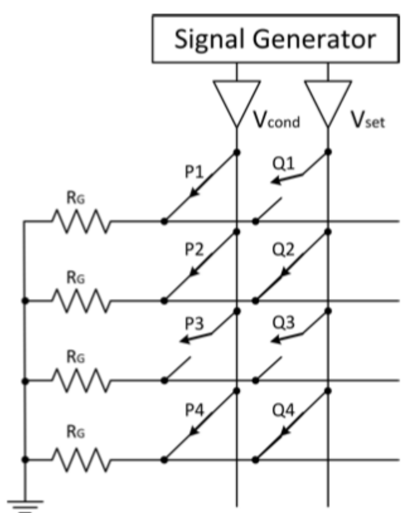

Designed versus Intrinsic Computation

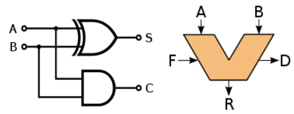

Desired computation

Computing architecture

Intrinsic dynamics

Device network

Adaptive programming

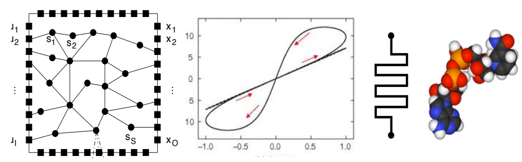

There is no general design theory on how to obtain a desired computation from intrinsic dynamics for devices that behave beyond simple Boolean switching.

LETTERS
PUBLISHED ONLINE: 21 SEPTEMBER 2015 | DOI: 10.1038/NNANO.2015.207
nature nanotechnology

Evolution of a designless nanoparticle network into reconfigurable Boolean logic

S. K. Bose[1†‡], C. P. Lawrence[1,2‡], Z. Liu[1], K. S. Makarenko[1], R. M. J. van Damme[3], H. J. Broersma[2] and W. G. van der Wiel[1*]

2-input logic gate
6 control variables

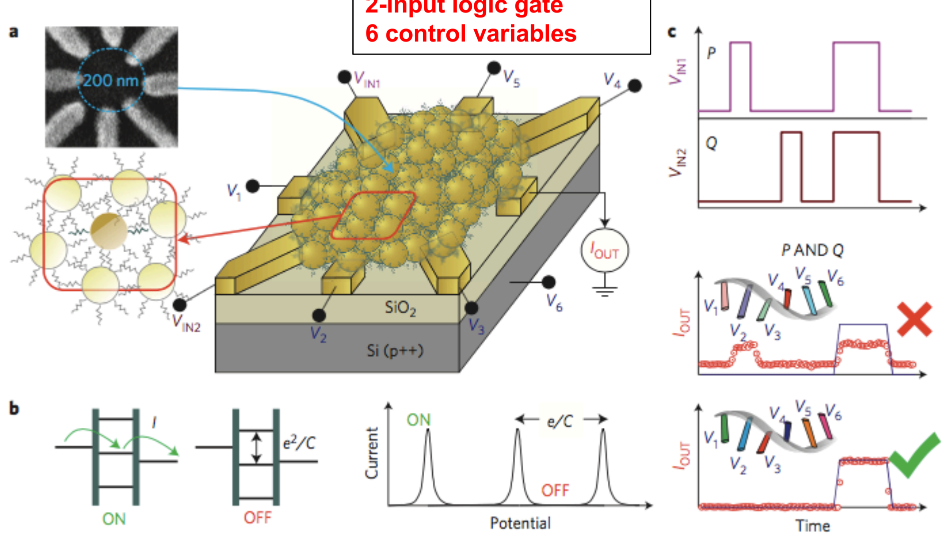

Reservoir Computing / Liquid State Machines

- Fixed reservoir with “interesting” dynamics. No state needed.
- Only the output layer is trained. → Low learning complexity.
- Variation is good! → Easy to fabricate.

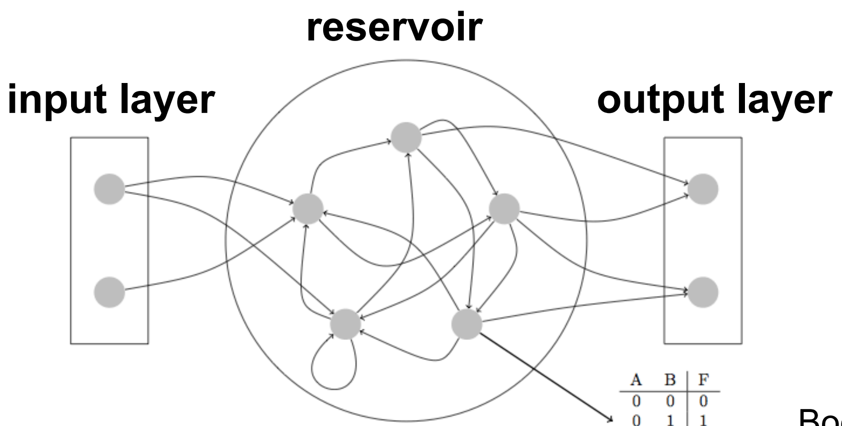

A	B	F
0	0	0
0	1	1
1	0	1
1	1	0

Boolean
or analog

Reservoir Computing / Liquid State Machines

[Obst et al., 2013] [Fernando and Sojakka, 2003]

Quantum dots

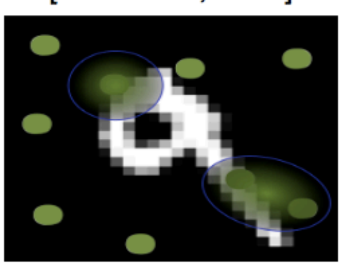

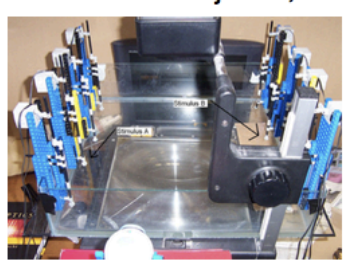

Water bucket

Atomic switches

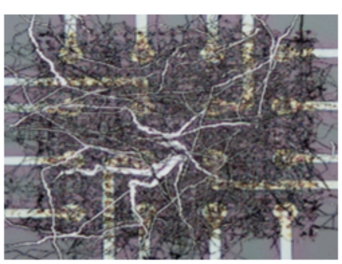

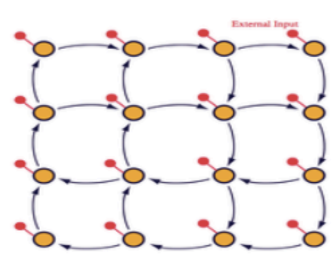

Photonics

[Sillin et al., 2013] [Vandoorne et al., 2011]

Reservoir Computing / Liquid State Machines

1 1 1 1
1 1 1 1

non-temporal

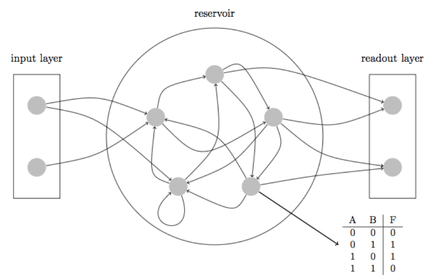

0 0 0 0
1 1 1 1

0 1 0 0
1 0 0 1

temporal

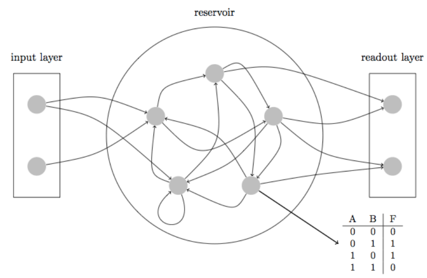

1 0 1 0
0 0 1 0

Issues and Challenges

Lack of composability

Lack of scalability

Monolithic systems
Signal attenuation

Solve large-scale real-world problems

Training

Hierarchical and modular systems

1. Create digital building blocks, then use traditional design tools?
2. New approach?

Hierarchical Networks

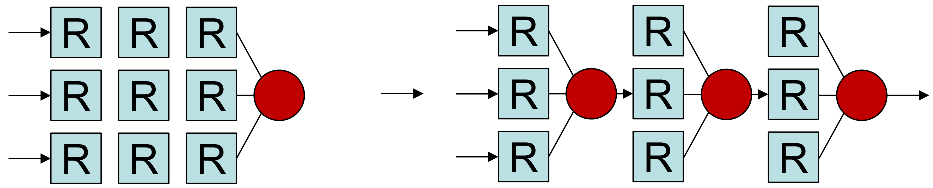

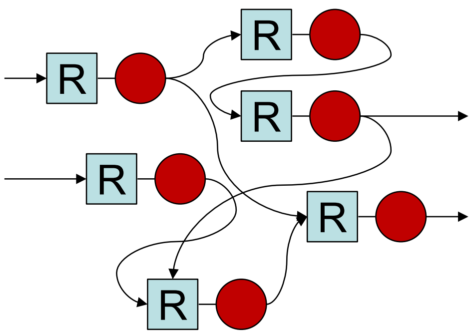

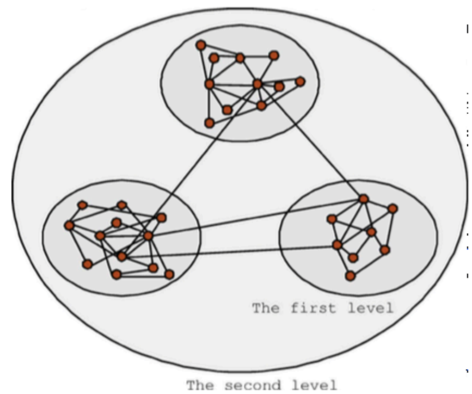

Hierarchical Composition of Memristive Networks for Real-Time Computing

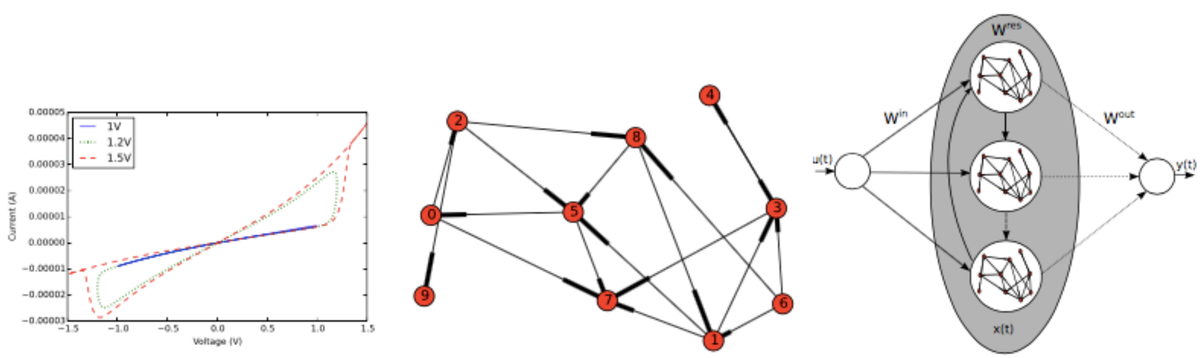

(a) Memristor switching (b) Memristive network (c) Memristive SCR

Burger et al, Nanoarch, 2015

Hierarchical Composition of Memristive Networks for Real-Time Computing

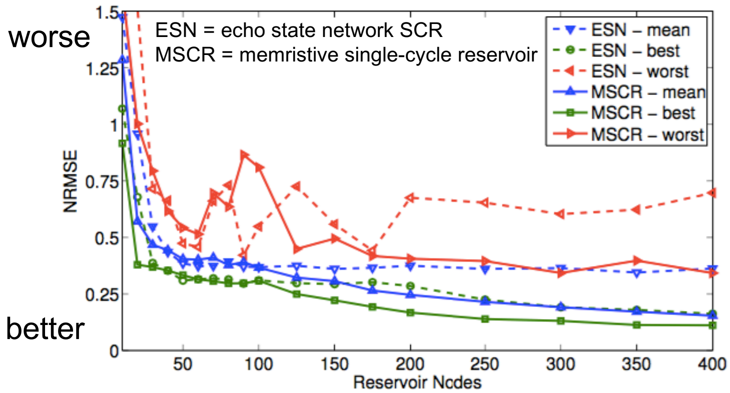

Fig. 4: NARMA performance for increasing SCR sizes. The dashed lines are obtained with regular sigmoidal neurons (ESN) and $\upsilon = 0.01$ and $\lambda = 0.75$. The memristive SCR (MSCR) data was obtained based on $\upsilon = 0.5$. $\lambda = 1.7$, and 52 memristive devices on average per node.

Burger et al, Nanoarch, 2015

- Hierarchical composition of heterogeneous small networks outperforms monolithic memristive networks by at least **20%** on waveform generation tasks.
- On the NARMA-10 task, we reduce the error by up to a factor of 2 compared to homogeneous reservoirs with sigmoidal neurons.
- Single memristive networks are unable to produce the correct result.

Network Topologies: Initial Steps

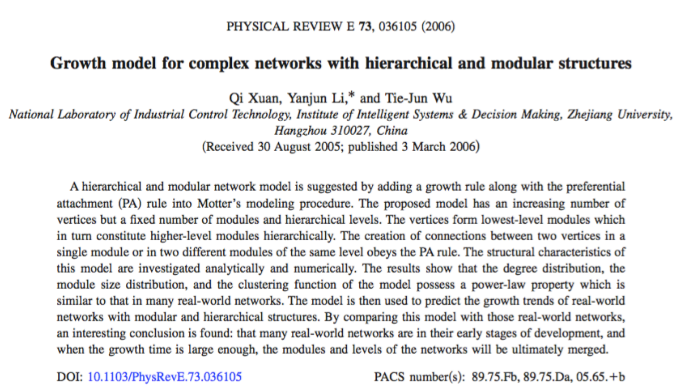

PHYSICAL REVIEW E **73**, 036105 (2006)

Growth model for complex networks with hierarchical and modular structures

Qi Xuan, Yanjun Li,* and Tie-Jun Wu

National Laboratory of Industrial Control Technology, Institute of Intelligent Systems & Decision Making, Zhejiang University, Hangzhou 310027, China

(Received 30 August 2005; published 3 March 2006)

A hierarchical and modular network model is suggested by adding a growth rule along with the preferential attachment (PA) rule into Motter's modeling procedure. The proposed model has an increasing number of vertices but a fixed number of modules and hierarchical levels. The vertices form lowest-level modules which in turn constitute higher-level modules hierarchically. The creation of connections between two vertices in a single module or in two different modules of the same level obeys the PA rule. The structural characteristics of this model are investigated analytically and numerically. The results show that the degree distribution, the module size distribution, and the clustering function of the model possess a power-law property which is similar to that in many real-world networks. The model is then used to predict the growth trends of real-world networks with modular and hierarchical structures. By comparing this model with those real-world networks, an interesting conclusion is found: that many real-world networks are in their early stages of development, and when the growth time is large enough, the modules and levels of the networks will be ultimately merged.

DOI: 10.1103/PhysRevE.73.036105 PACS number(s): 89.75.Fb, 89.75.Da, 05.65.+b

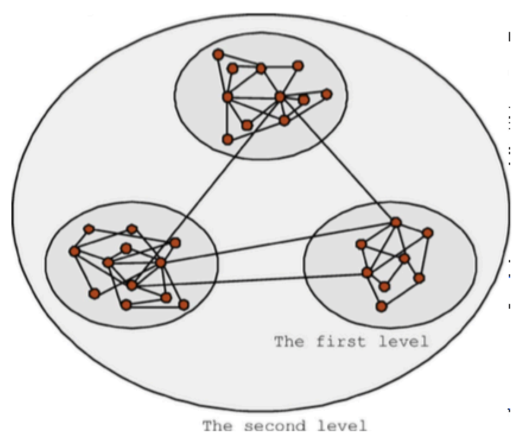

FIG. 2. (Color online) The network model with $M=2$, $n=3$, $m=m_0=2$, $q_1=0.9$, and growth time $T=30$. It is grown up to $N=31$ vertices.

- **M**: number of hierarchical levels
- **n**: number of modules grouped together to constitute the modules of the next hierarchical level

Random Boolean Network Reservoir

NK Networks:

- ***N*** = number of nodes
- ***K*** = interaction between the nodes, i.e., the number of incoming links per node

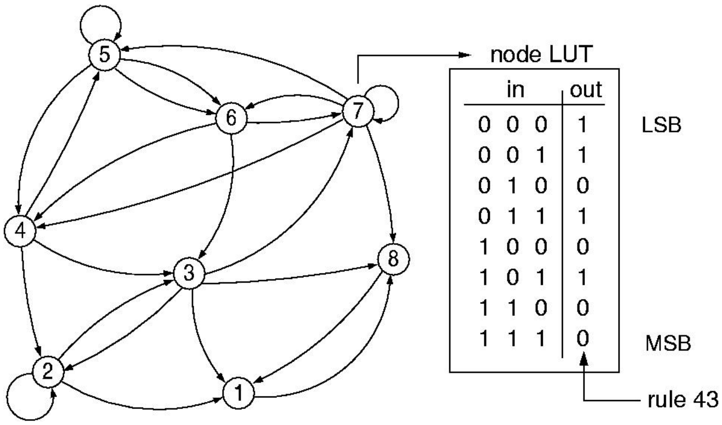

N=8, K=3

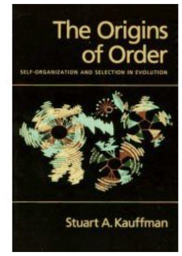

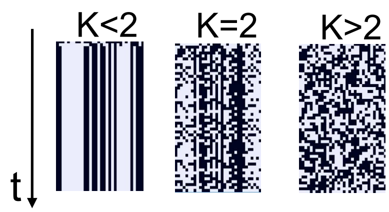

X-O Tasks

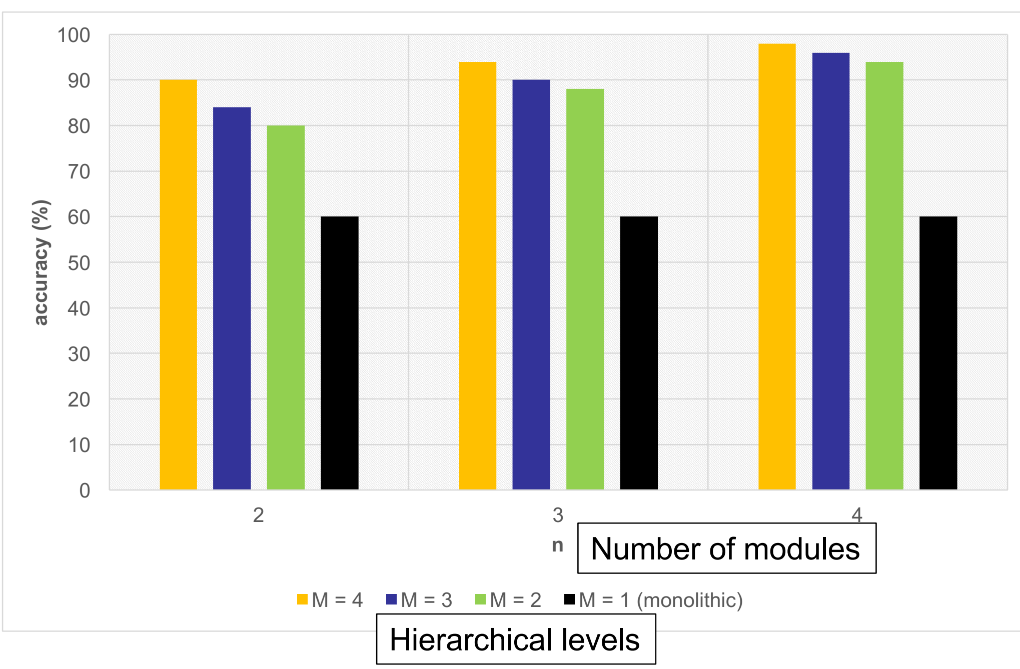

Labyrinth Tasks

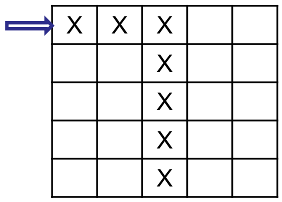

Fig. 6.1. 1-turn labyrinth

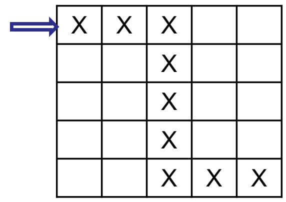

Fig. 6.2. 1-turn labyrinth

Labyrinth Tasks

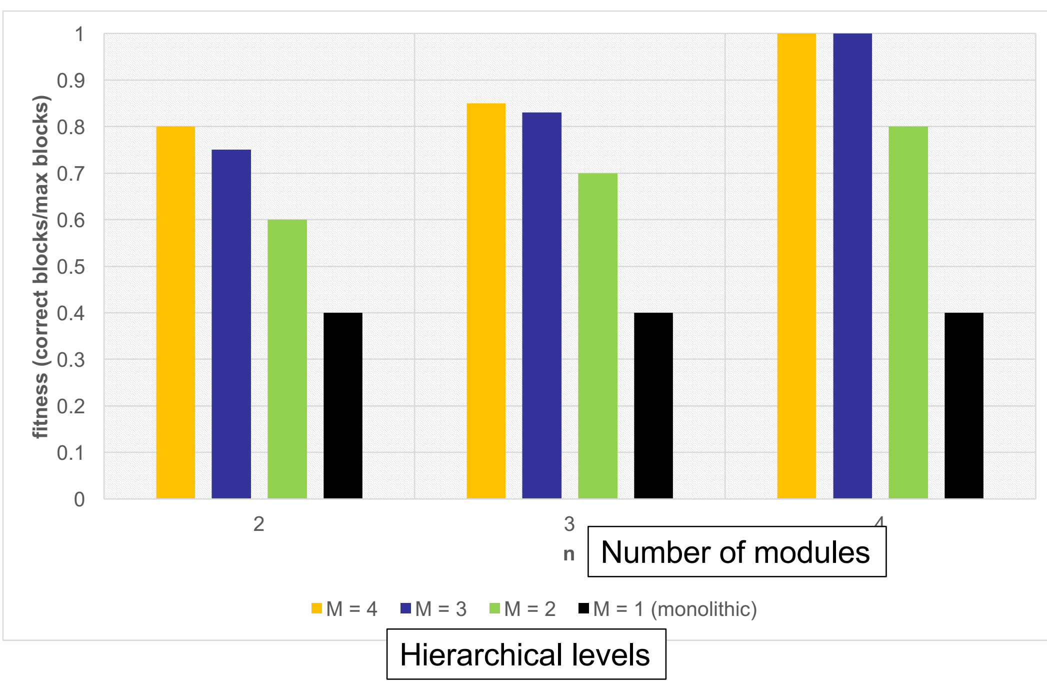

Deep LSM network (D-LSM)

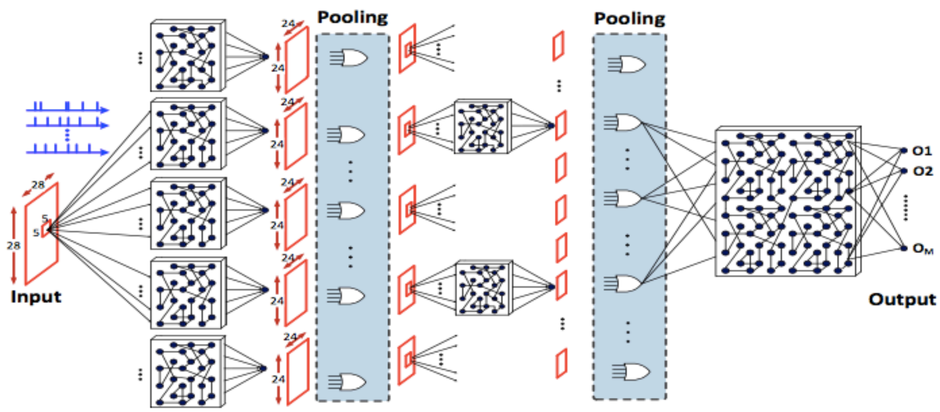

Fig. 4: The proposed deep LSM network. The image pixels are first converted to spike trains. Each 1st-stage LSM receives 25 spike trains which corresponds a sub-region of the input image selected by a 5×5 sliding window. As the sliding window moves to cover the entire image, each 1st-stage LSM generates 24×24 new spike trains. The pooling/sub-sampling stage is realized by 4-inputs OR gates. After multiple LSM stages and pooling stages, the extracted features enter the last LSM stage which also incorporates the final readout layer.

Wang and Li, **D-LSM: Deep Liquid State Machine with Unsupervised Recurrent Reservoir Tuning**, 3rd International Conference on Pattern Recognition (ICPR), December 4-8, 2016

Memcapacitive RC

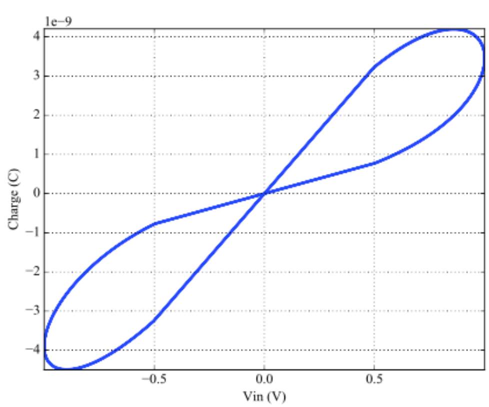

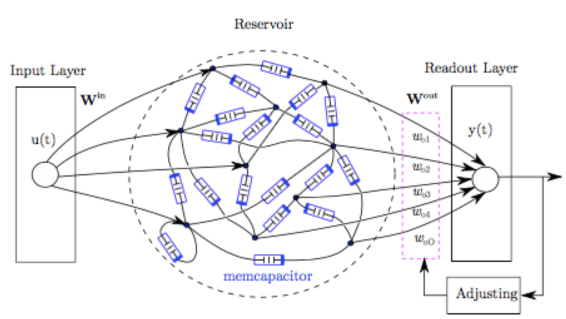

Memcapacitive RC: MNIST

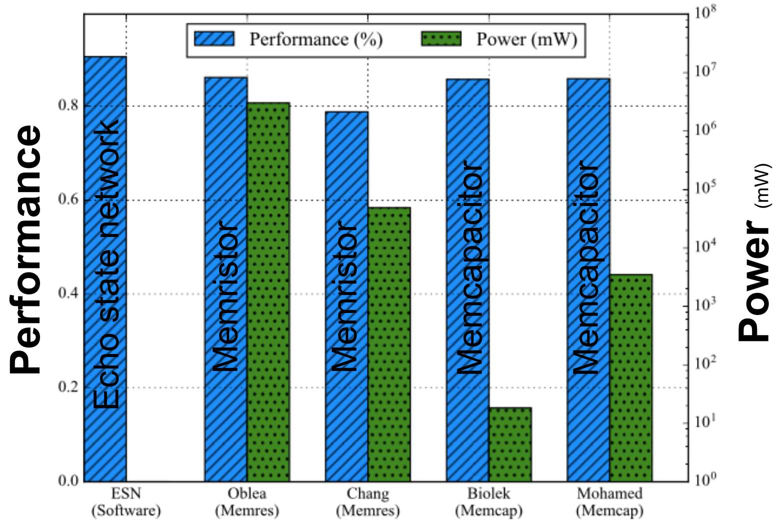

Fig. 4: Reservoir performance and power consumption for the MNIST task. $I = 784, O = 600, N = 2100$. The power measurements represent the average per image.

Thanks to Students and Sponsors

NSF

DARPA

SRC® Semiconductor Research Corporation

- SRC Education Alliance, Undergraduate Research Opportunities (URO) Program. Sep 2012 – Sep 2016.
- **Molecular Computing for the Real World**, National Science Foundation (NSF), CISE, SHF. NSF grant no: 1518833, Sep 1, 2015 – Aug 31, 2020. $209,545.
- **Unified English Braille through a Powerful and Responsive eLearning Platform (UEB PREP)**, Rehabilitation Services Administration, Department of Education. Sep 1, 2014 – Aug 31, 2019. $548,483
- DARPA, **Sparse Adaptive Local Learning for Sensing and Analytics (SALLSA)**. May 3, 2013 – Aug 2, 2017. The project is in collaboration with the University of Michigan and Los Alamos National Laboratory. $5.69 million.
- **Inference at the Nanoscale**, National Science Foundation (NSF), Cyber-enabled Discovery and Innovation (CDI), Type II award, NSF grant no: 1028378, Sep 15, 2010 – Aug 31, 2016 (with NCE)