A Probabilistic Self-Annealing Compute Fabric Based on 560 Hexagonally Coupled Ring Oscillators for Solving Combinatorial Optimization Problems

Ibrahim Ahmed[1], Po-Wei Chiu[1], and Chris H. Kim[1]

[1] University of Minnesota

Outline

- Motivation and Prior Art
- Circuit and Architecture Design
- Mapping Graph Problems to Hardware
- Measured Results
- Conclusion

Outline

- Motivation and Prior Art
- Circuit and Architecture Design
- Mapping Graph Problems to Hardware
- Measured Results
- Conclusion

Motivation

Autonomous vehicles

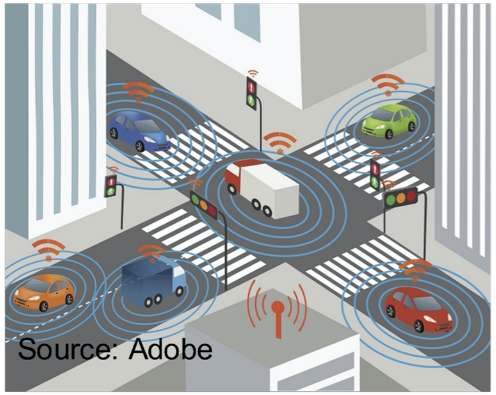

Communication networks

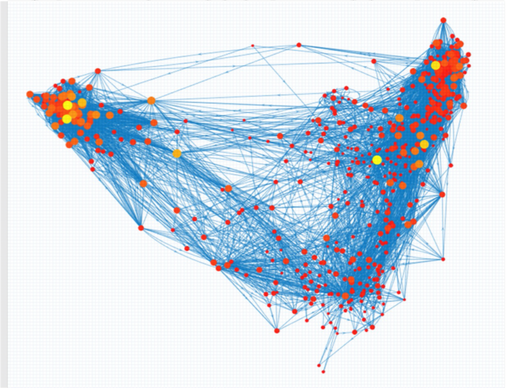

Smart grid

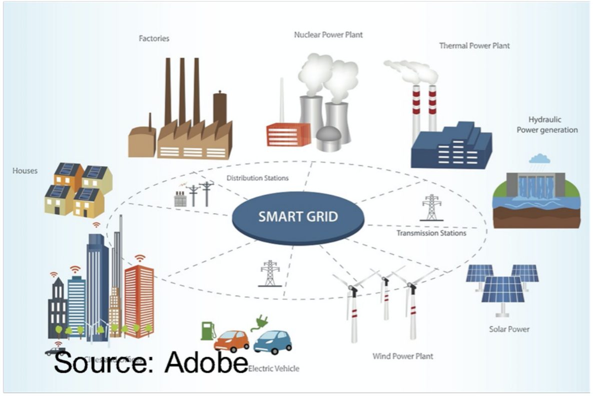

VLSI routing

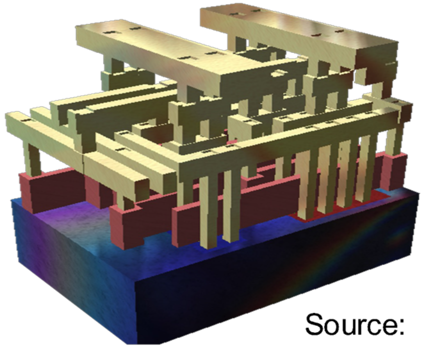

Source: Wikipedia

Applications of combinatorial optimization problems (COPs)

COP: Boolean satisfiability

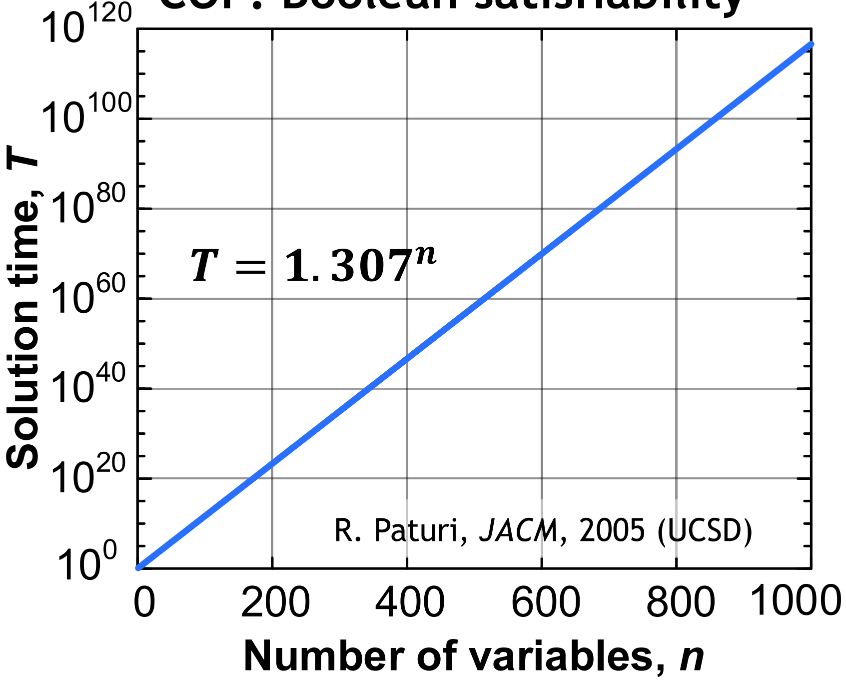

- Solution space and time increases rapidly with the number of variables
- Intractable to solve using traditional von Neumann computer

Motivation

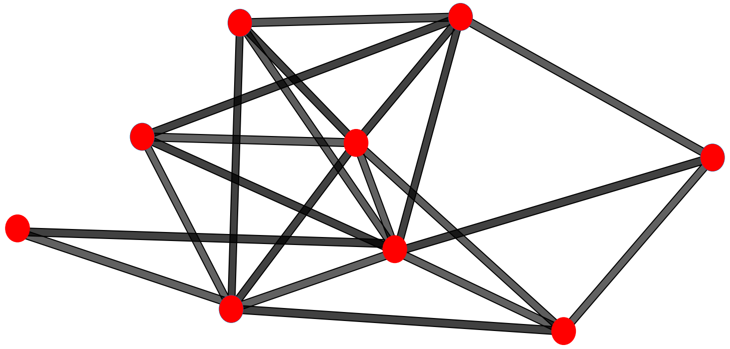

COP mapped to graph

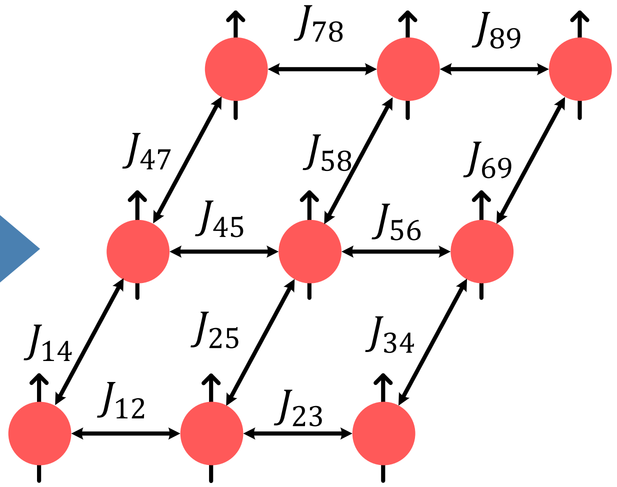

Network of Spin

$$H(s) = -\sum_{i,j} J_{ij} \cdot s_i s_j - \sum_i h_i s_i$$

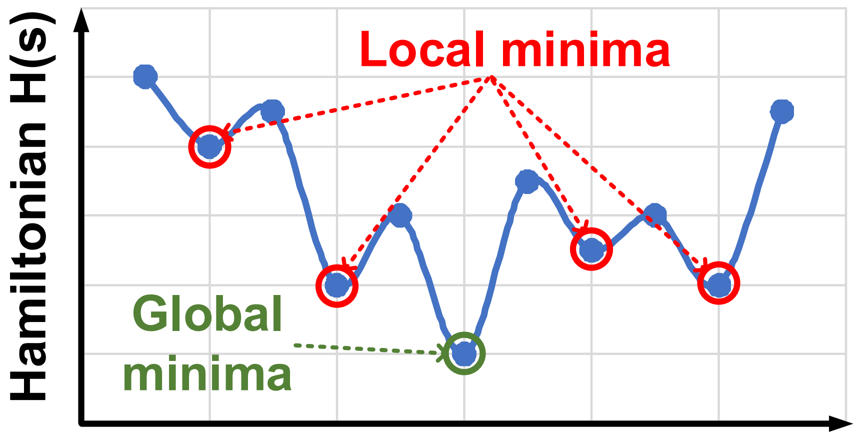

Spin states

J_{ij} = Coupling between spins i and j
s_i = State of spin i, h_i = Local field to spin i
$H(s)$ = Hamiltonian of the system

- Ising computation uses energy transfer between spins to reach global minima
- A network of coupled oscillators can solve many NP-hard and NP-complete problems, such as COPs

A. Lucas, *Front. Phys.*, 2014

Prior Art

Quantum computer

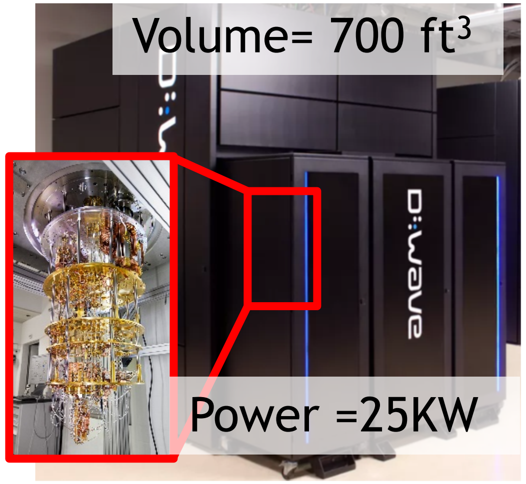

D-Wave

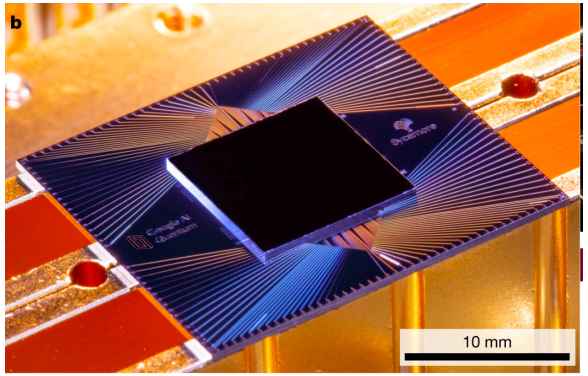

Google

Pros:

Can theoretically solve a wide variety or problems

Cons:

1. Requires sophisticated hardware
2. Very sensitive to noise
3. Operating temperature -273.14°C

Special process

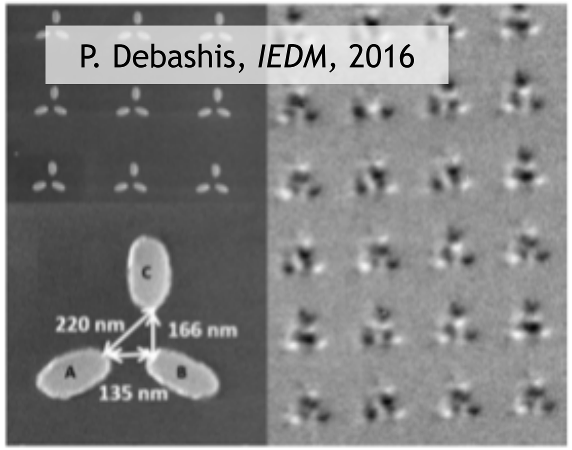

Nano-magnet based

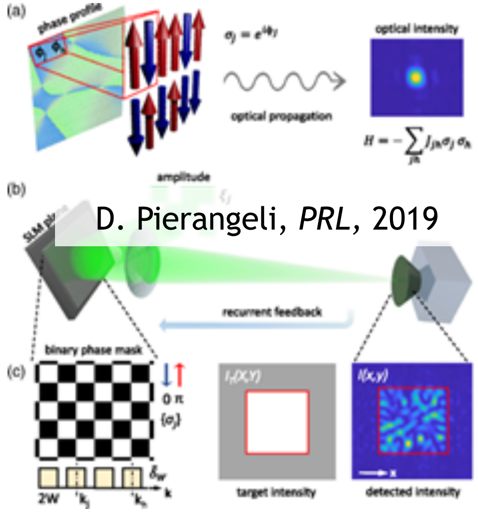

LASER based

Pros:

Can achieve coupling dynamics

Cons:

1. Requires special process with fabrication challenges
2. Limited practical demonstrations

Prior Art

ASIC implementation

M. Yamaoka, *JSSC*, 2016
T. Takemoto, *ISSCC*, 2019

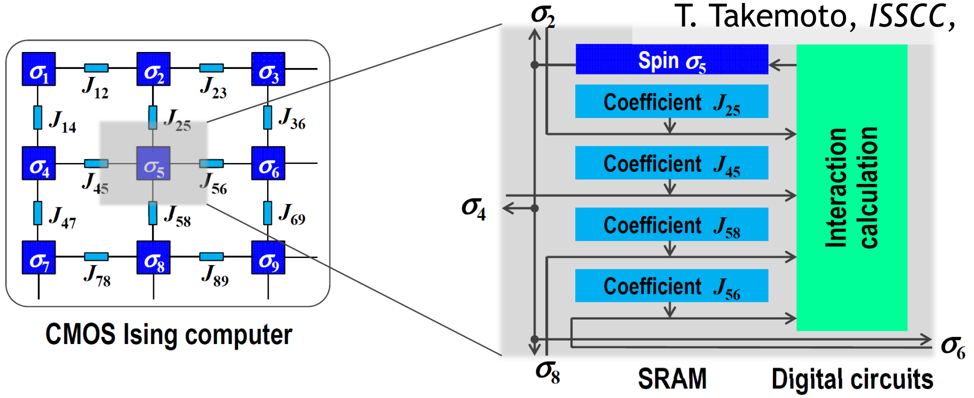

Pros:
Can theoretically solve a wide variety or problems

Cons:
1. No coupling dynamics : computation time and energy may be higher
2. Search for better solution is time dependent

Breadboard and PCB implementation

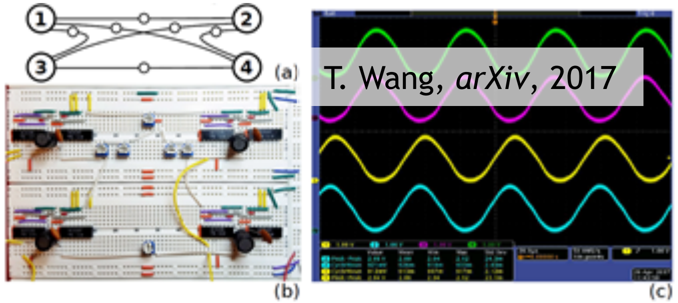

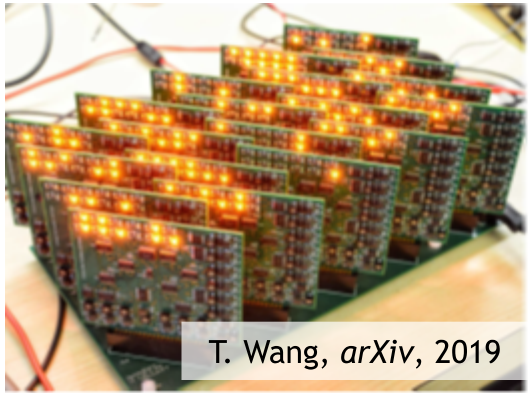

Pros:
Can achieve coupling dynamics

Cons:
Not practical for large number of programmable spins

CMOS integrated Ising computer has the best of both worlds

Outline

- Motivation and Prior Art
- Circuit and Architecture Design
- Mapping Graph Problems to Hardware
- Measured Results
- Conclusion

Choice of CMOS Spin

LC oscillator:
- Pros: low phase noise, frequency insensitive to PVT
- Cons: very large area, narrow frequency tuning range

Ring oscillator (ROSC):
- Pros: small area, simple design, wide frequency tuning range
- Cons: higher phase noise, higher sensitivity to PVT

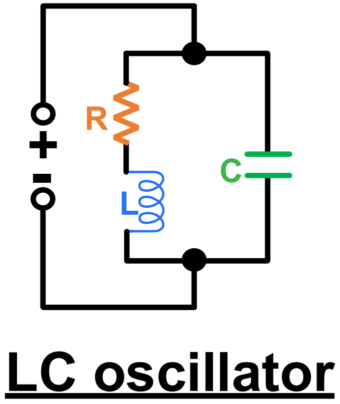

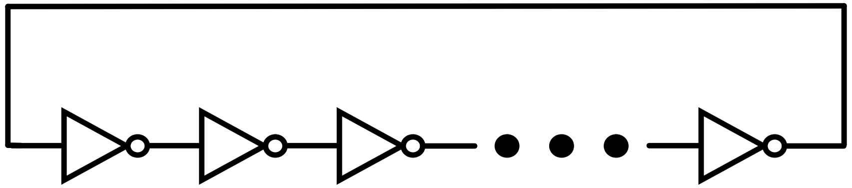

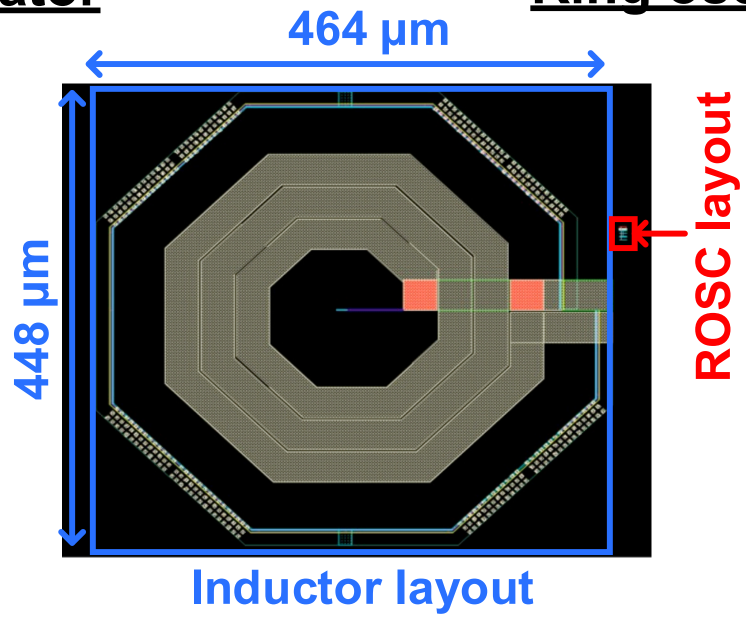

Type of oscillator not important for solving Ising problems (e.g. mechanical, electrical, chemical, etc.)

ROSC Coupled Using Digital Latches

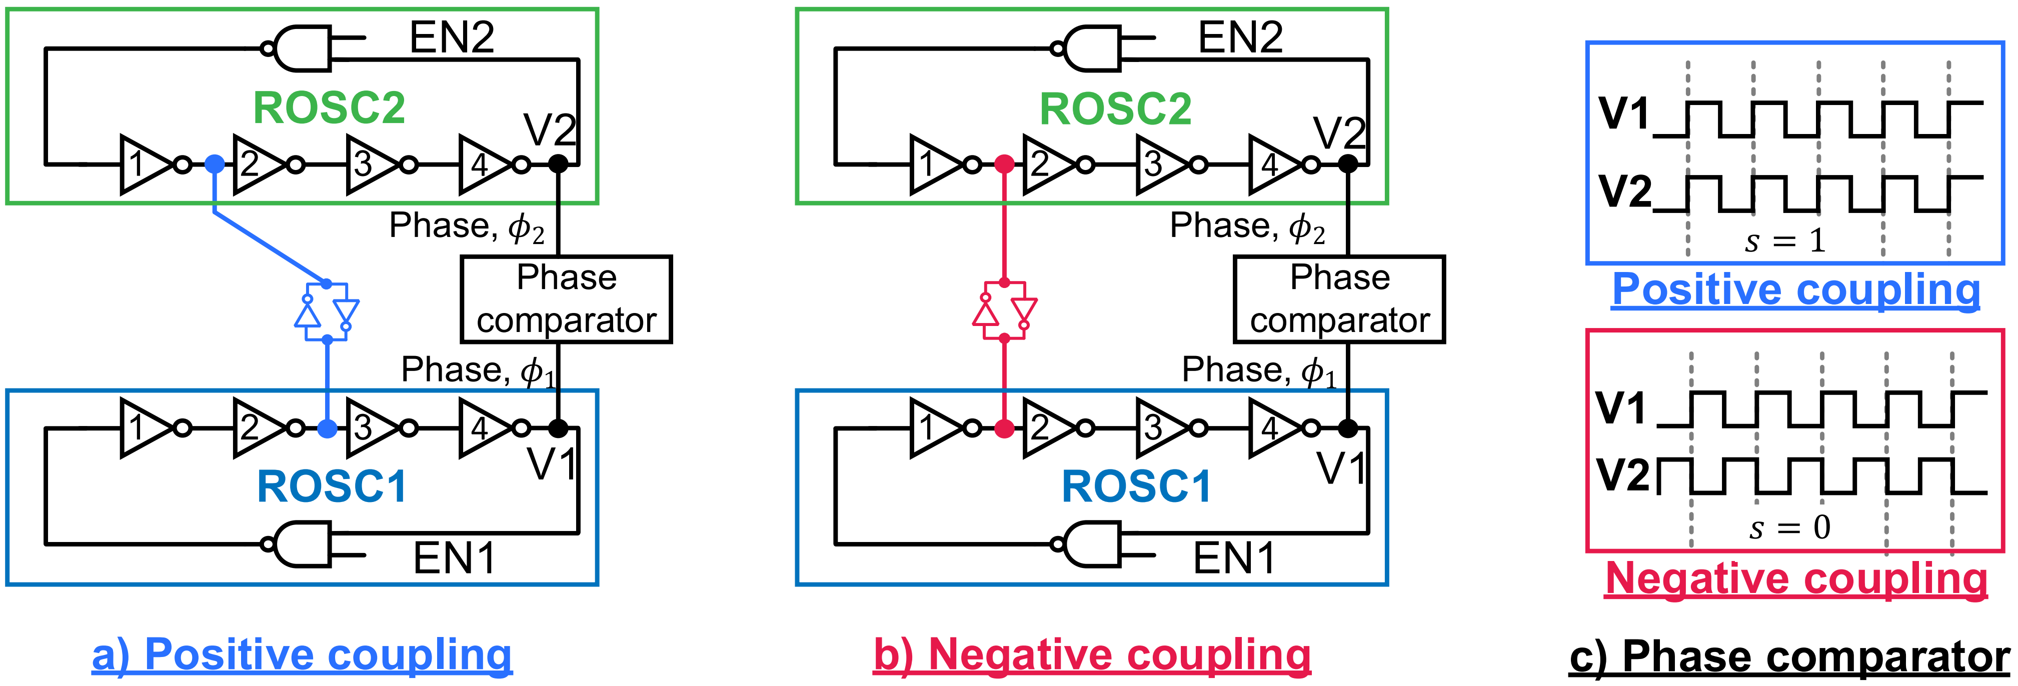

- Any coupling medium that enables energy transfer may couple ROSCs
- ROSC and digital latches are designed with global and local enable signals

Choice of Architecture

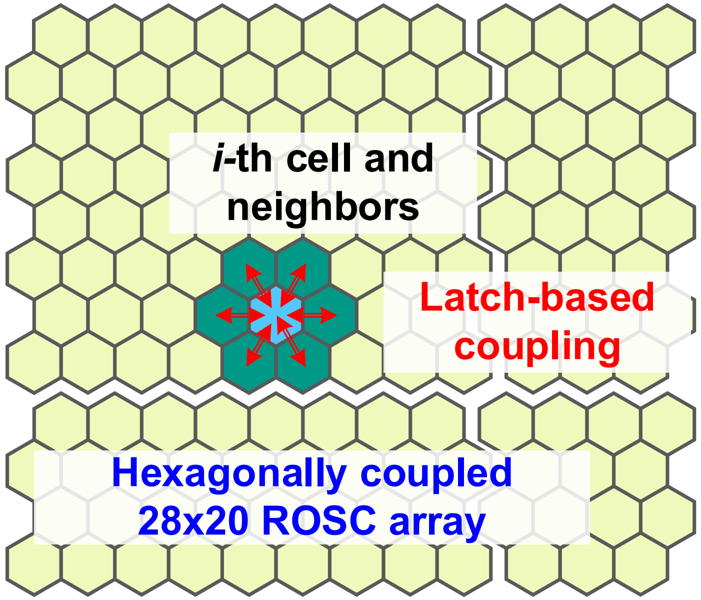

<u>Proposed architecture</u>

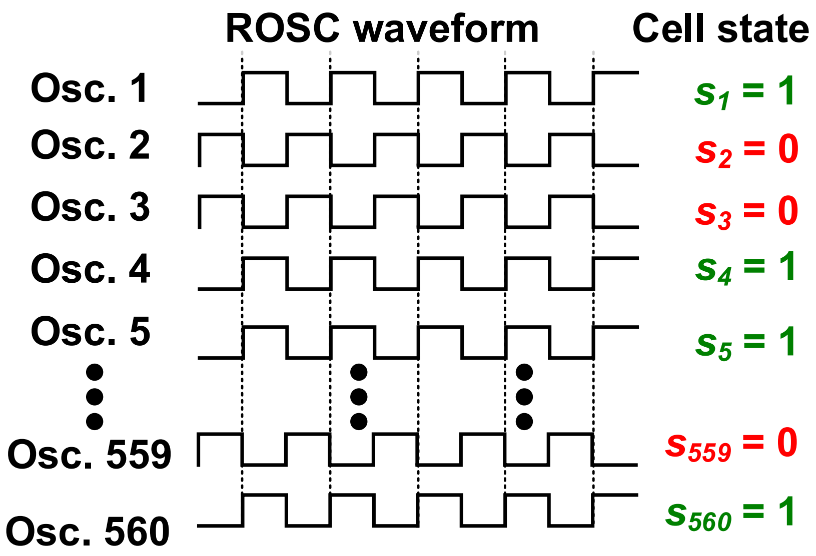

<u>ROSC waveform and unit cell states</u>

$$E_{system} = -\sum_{i=1}^{N}\sum_{j=1}^{n_i} J_{ij} \cdot s_i \cdot s_j$$

n_i = number of coupled neighbors of i-th cell,
J_{ij} = coupling weight between cells i and j

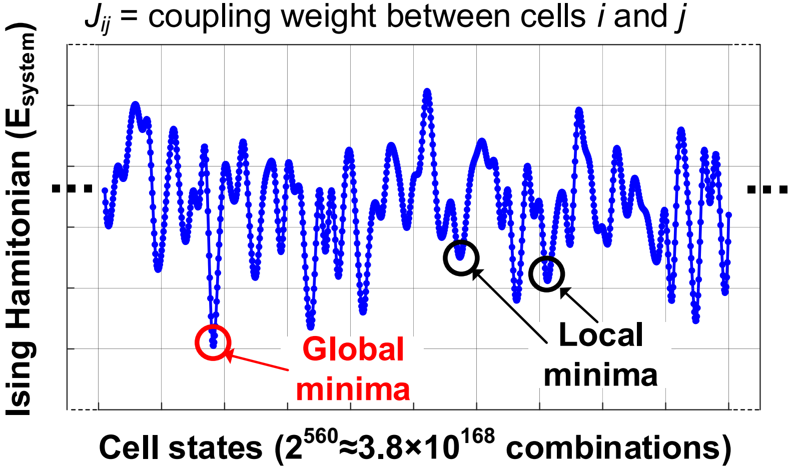

<u>System energy of the Ising computer</u>

- Hexagonal unit cell maximizes the number of neighbors in 2D plane
- Latch based coupling between cells is digitally controlled

ROSC and Latch Design

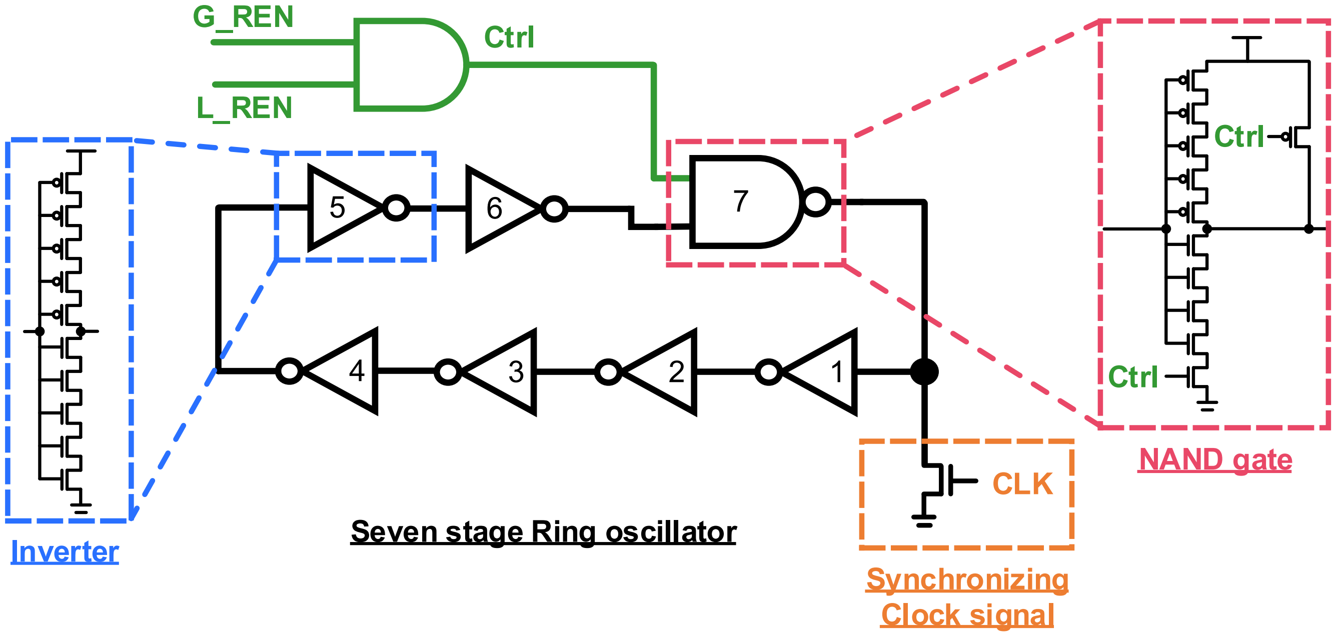

ROSC:

- Seven stage ROSC
- Global and local enable signals, G_REN and L_REN, respectively, controls ROSC activation
- Measured frequency: 120MHz

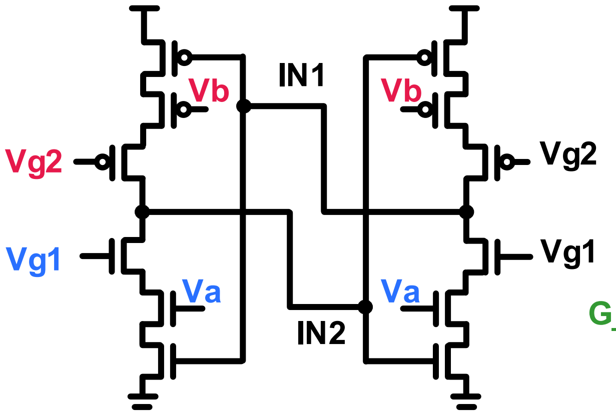

Digital latch

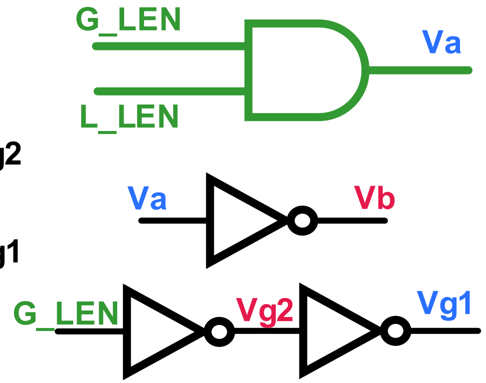

Latch Control signals

Digital Latch:

- Designed with back to back inverters
- Global and local enable signals, G_LEN and L_LEN, respectively, controls latch activation

Modular Unit Cell: Circuit Blocks

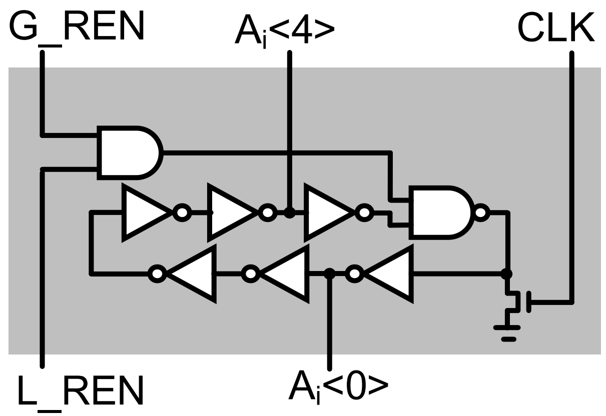

a) ROSC block

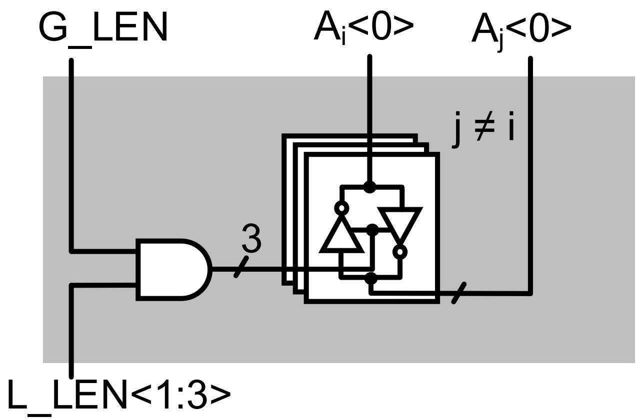

b) Latch coupling block

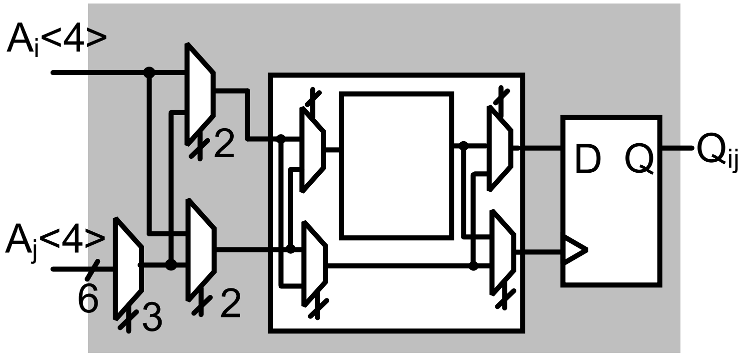

c) Read block

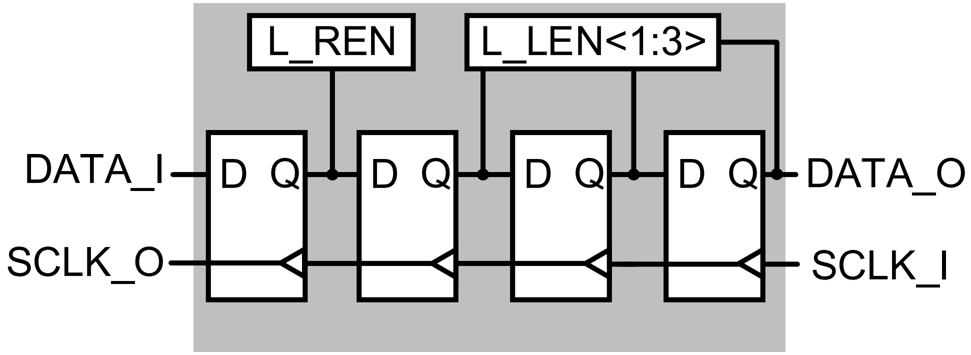

d) Scan block

- Read block samples one of the neighboring ROSC
- Scan block programs the graph: four program bit per cell, 2240 bit for the chip

Timing Based Annealing Mechanism

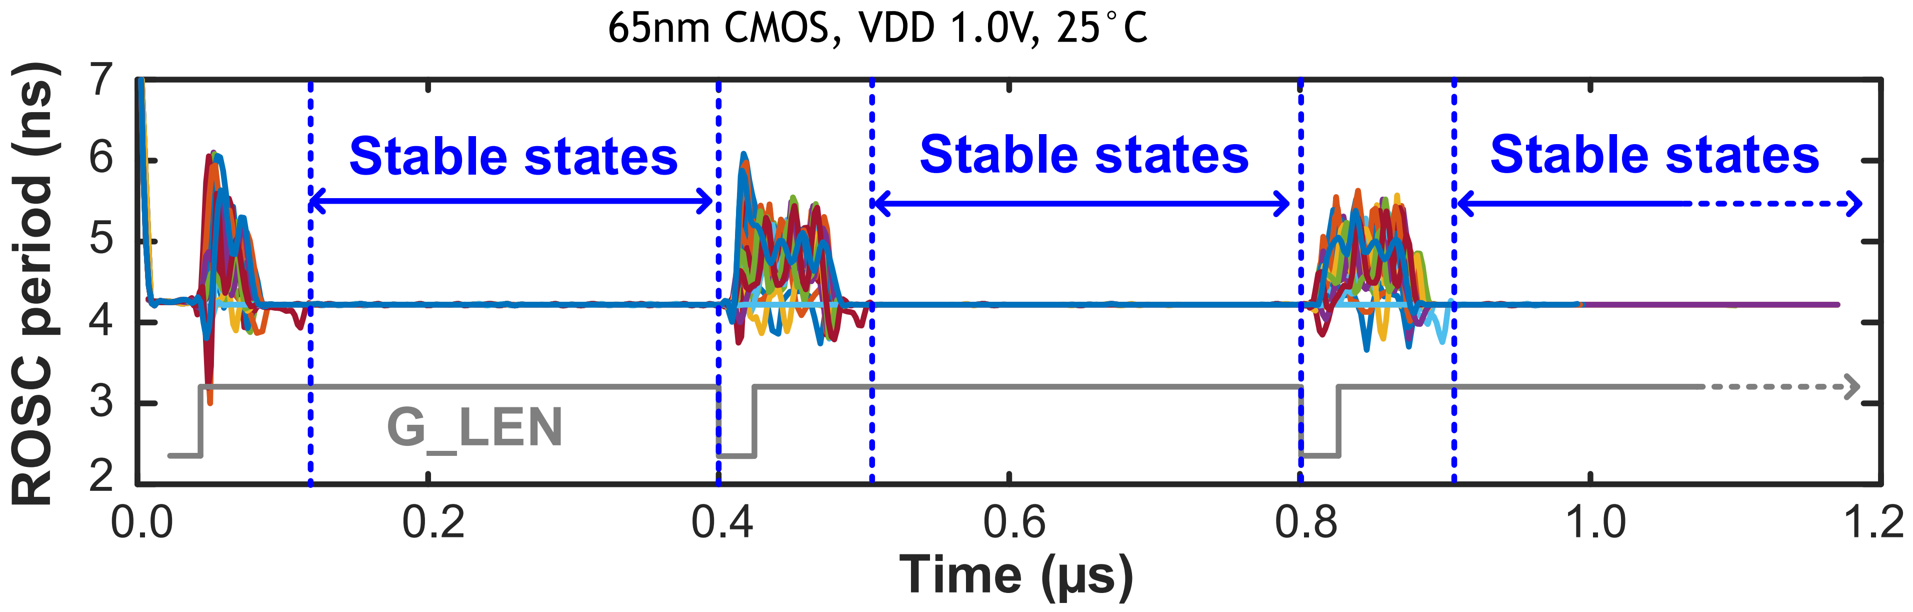

- Annealing helps the chip to find a better local minima
- We periodically turn off the coupling latches and let spin states to change
- To reduce delay, we only annealed the chip three times for measurement

Die Photo and Chip Summary

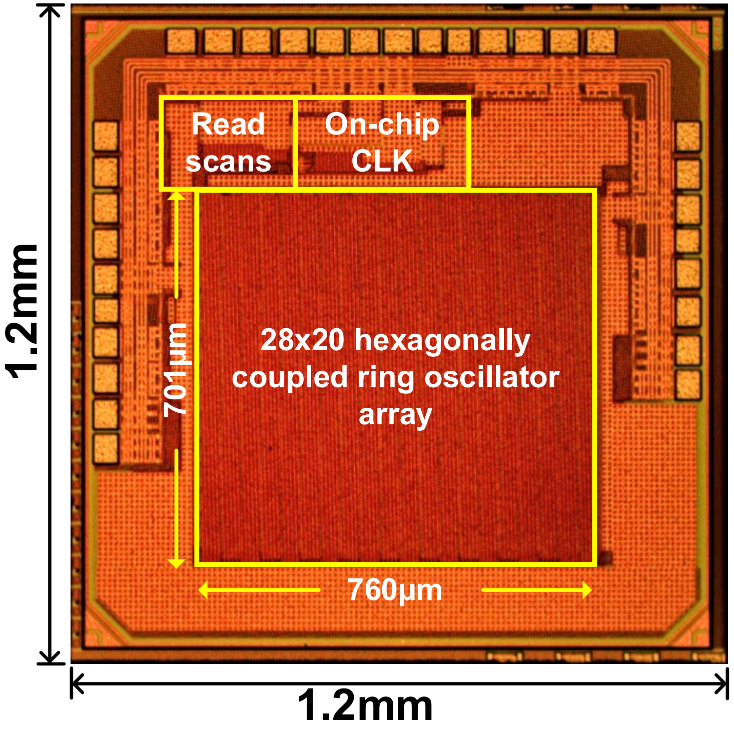

Die photo

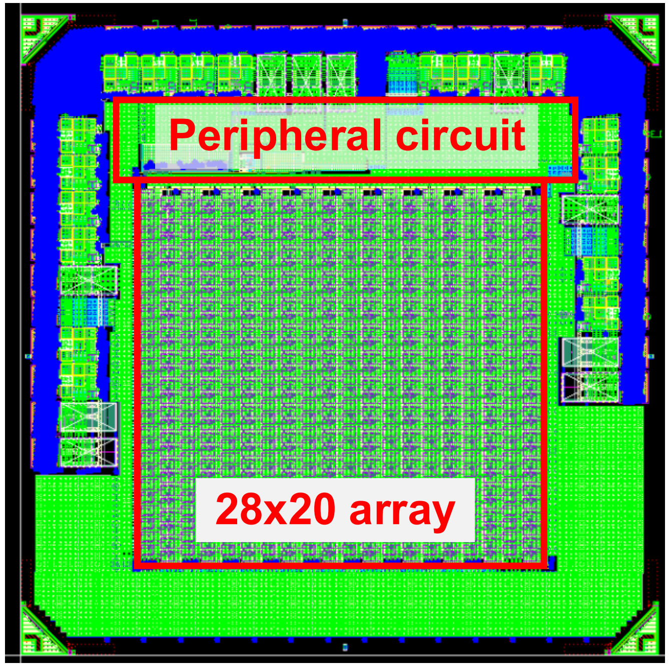

Full chip layout

Application	Combinatorial optimization problems
Process	65nm CMOS
Architecture	ROSC, latch based coupling, self-annealing
Voltage	1.0V
Area	Chip: 1.44mm^2
	Core: 0.53mm^2
	Unit cell: 0.00095mm^2
Peak power	23mW
Power per cell	41μW

Chip summary

- 28x20=560 coupled oscillators (only limiting factor is chip area)

Outline

- Motivation and Prior Art
- Circuit and Architecture Design
- Mapping Graph Problems to Hardware
- Measured Results
- Conclusion

Mapping Graph Problems to Hardware

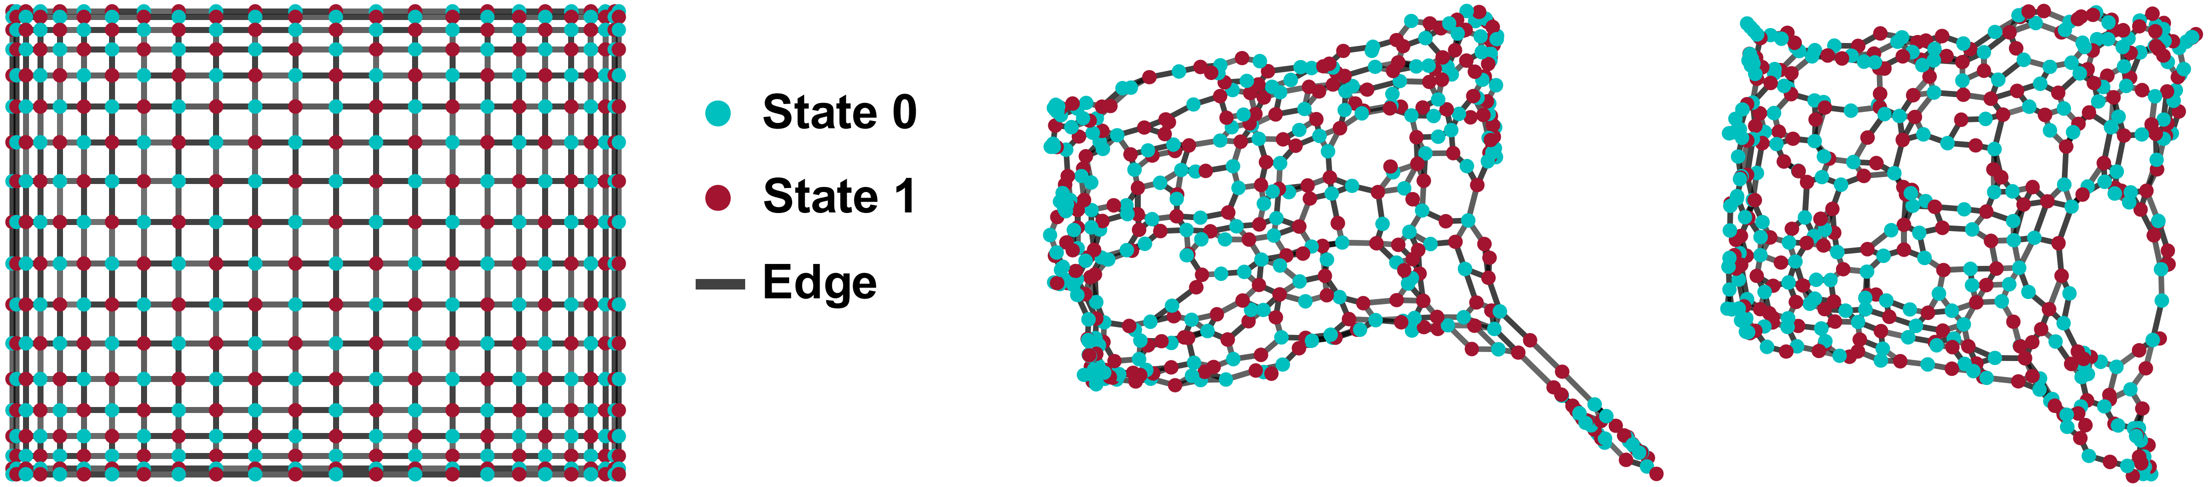

a) "Easy" COP (measured)

b) "Difficult" COP (measured)

- Randomly generate spatial graph using an algorithm that can be directly mapped to the chip without complicated embedding algorithm
- "Easy" COP: 2-color graphs such as checker board pattern
- "Difficult" COP: 4 or more color graphs

Introduction to Max Cut Problem

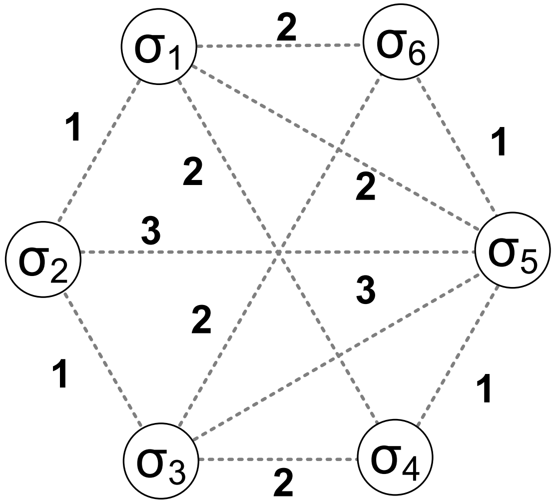

- The problem of finding a maximum cut in a graph is known as the Max-Cut Problem
- Finding max-cut of a graph is an *NP-hard* problem

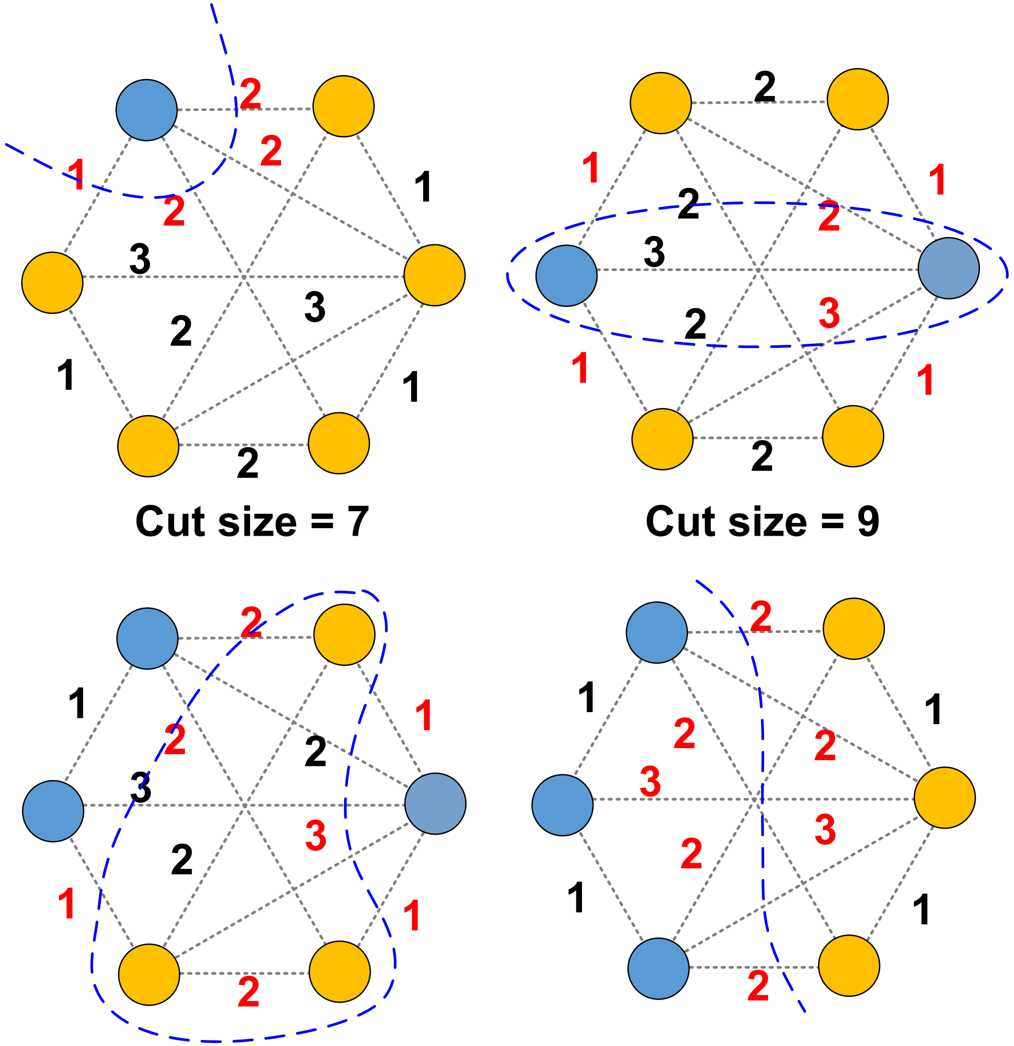

Cut size = 7

Cut size = 9

Cut size = 12

Cut size = 16 (max)

Introduction to Max Cut Problem

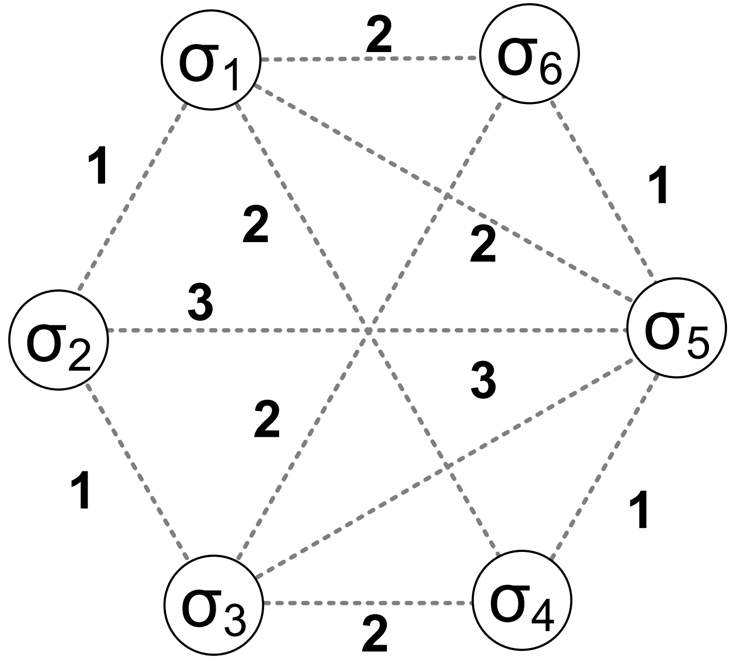

$$H(\sigma) = -\sum_{i,j}(-w_{ij})\,\sigma_i\sigma_j$$
$$= \sum_{diff\ group}(-w_{ij}) + \sum_{same\ group} w_{ij}$$
$$= \sum_{diff\ group}(-w_{ij}) + \left[\sum_{i,j} w_{ij} - \sum_{diff\ group} w_{ij}\right]$$
$$= \sum_{all} w_{ij} - 2\times \sum_{diff\ group} w_{ij}$$

Cut size

H = Hamiltonian of the system
σ_i = Spin status of magnet i {+1 or -1}
w_{ij} = weight between magnets i and j

- Ising Hamiltonian = [sum of all weights] – 2×[cut size]

Outline

- Motivation and Prior Art
- Circuit and Architecture Design
- Mapping Graph Problems to Hardware
- Measured Results
- Conclusion

Repeated Measurement of Same Graph

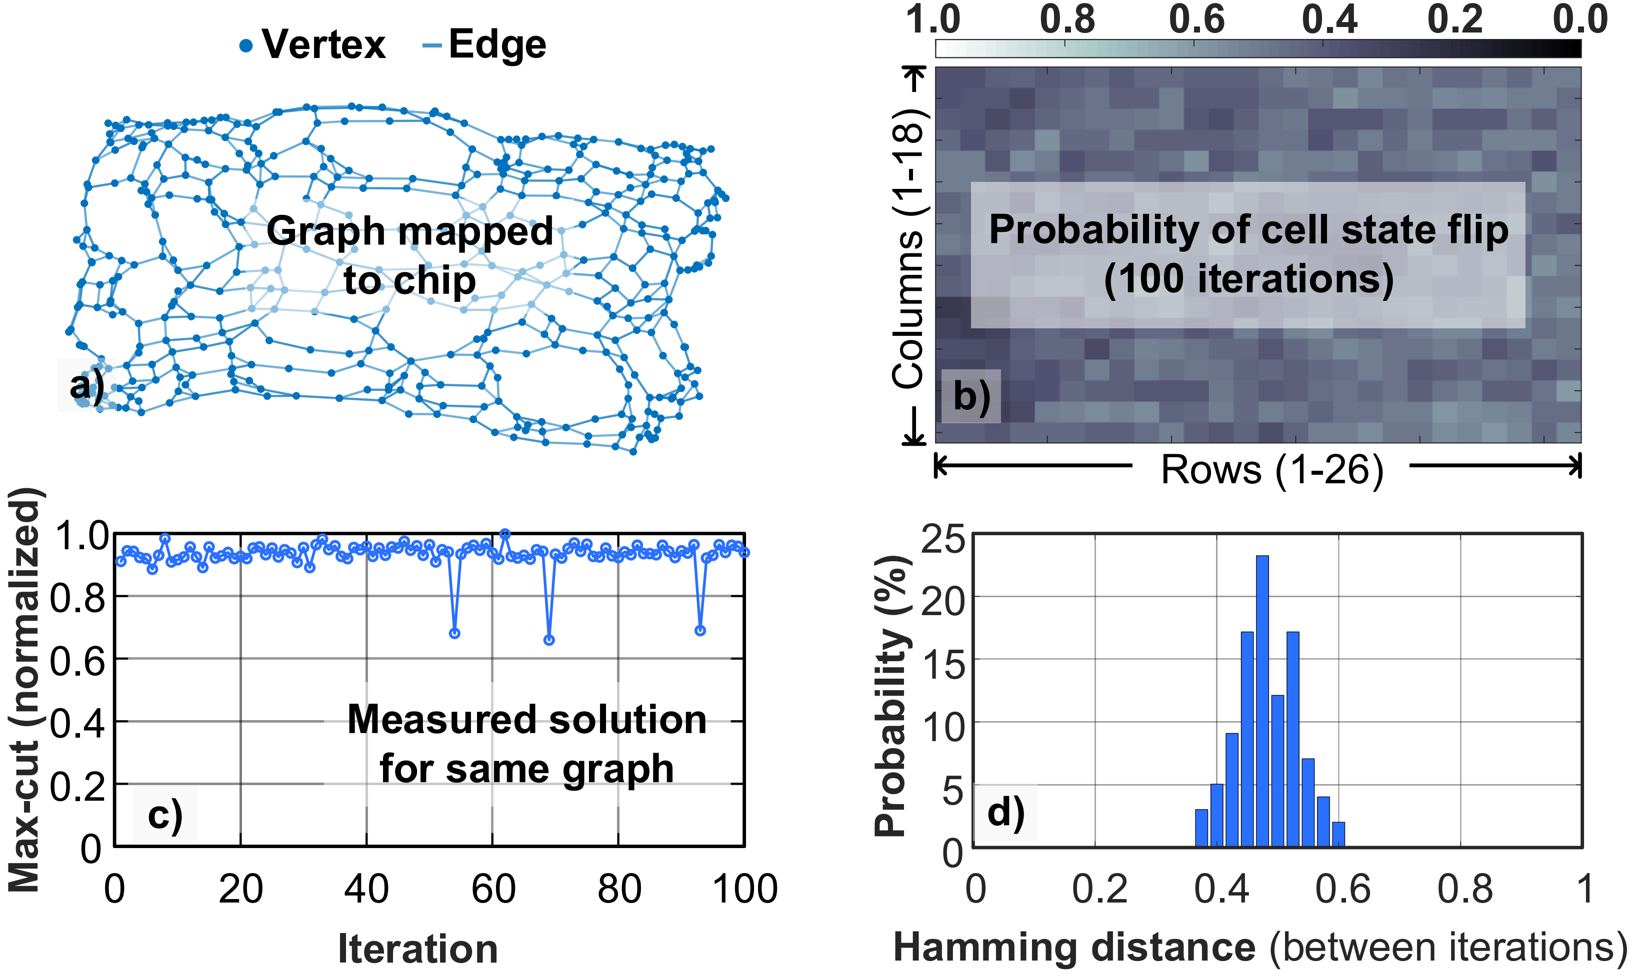

- Unit cells may change states between iterations while the final cut values remain similar : proving chip is **exploring various local minima**

Max-Cut Results for Graph Size 15x15

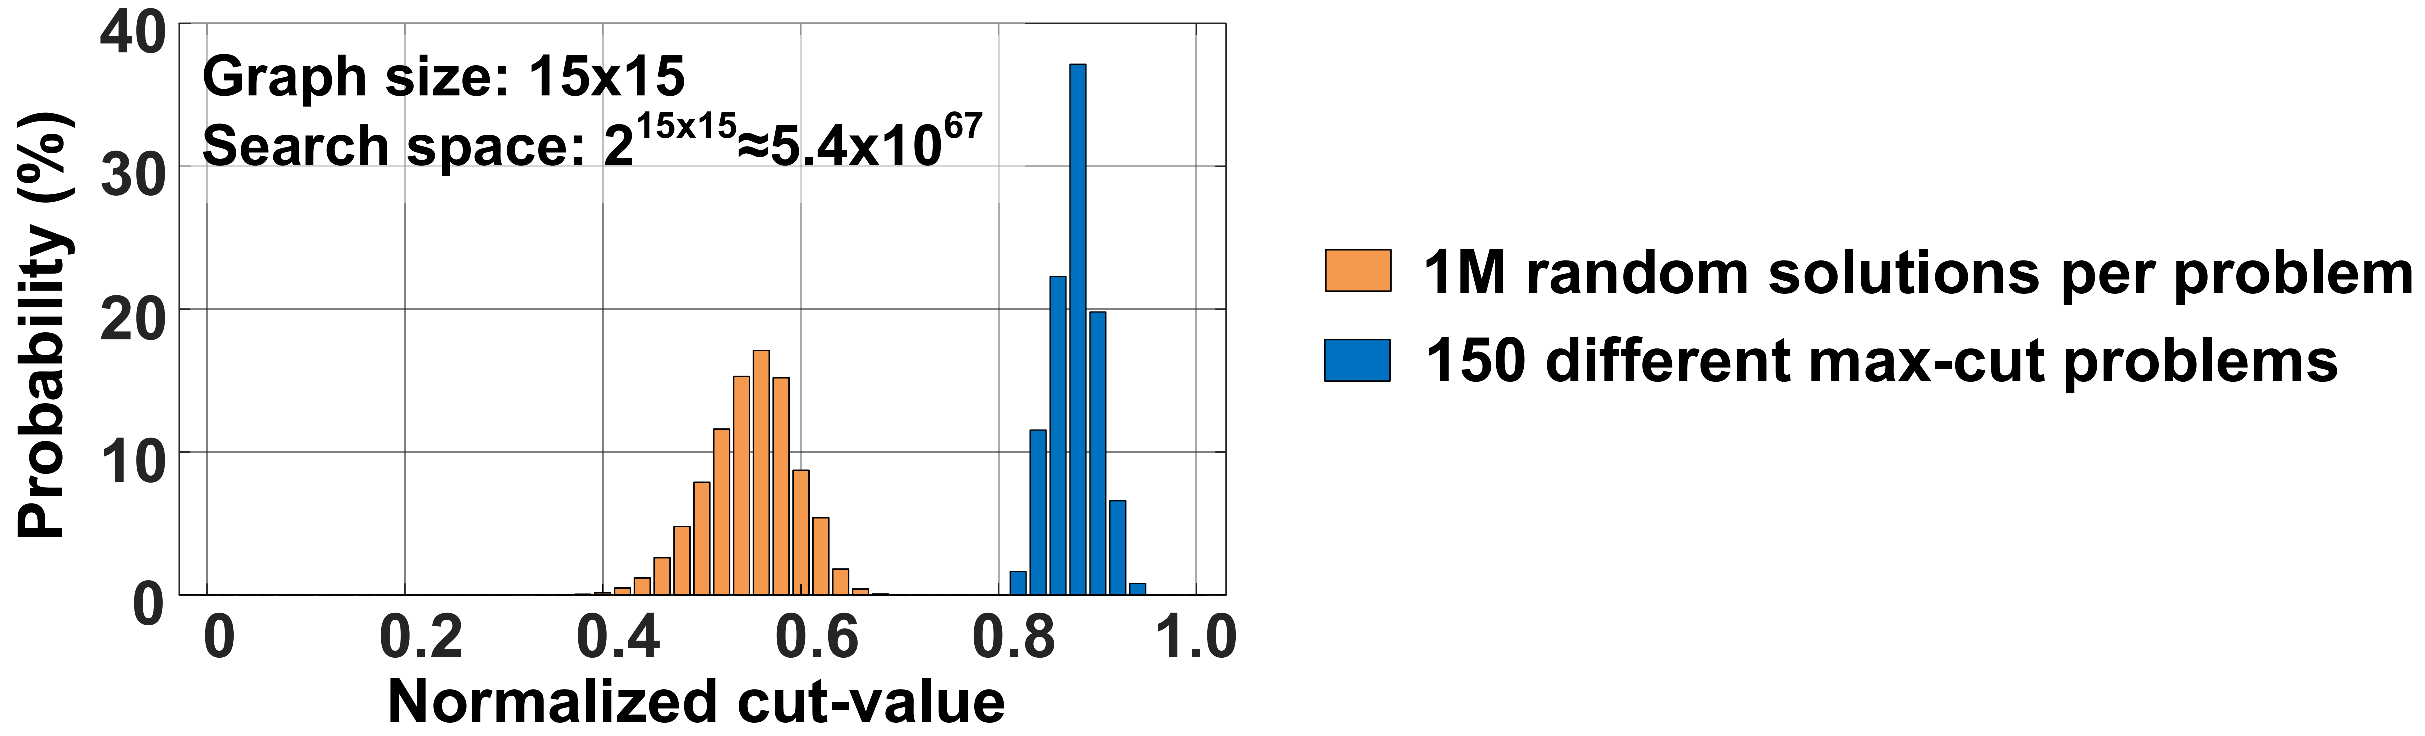

- 150 difficult COPs are mapped and max-cut results are measured
- Annealing is done only three times as compared to thousands of times in quantum computers

Max-Cut Results for Different Graph Sizes

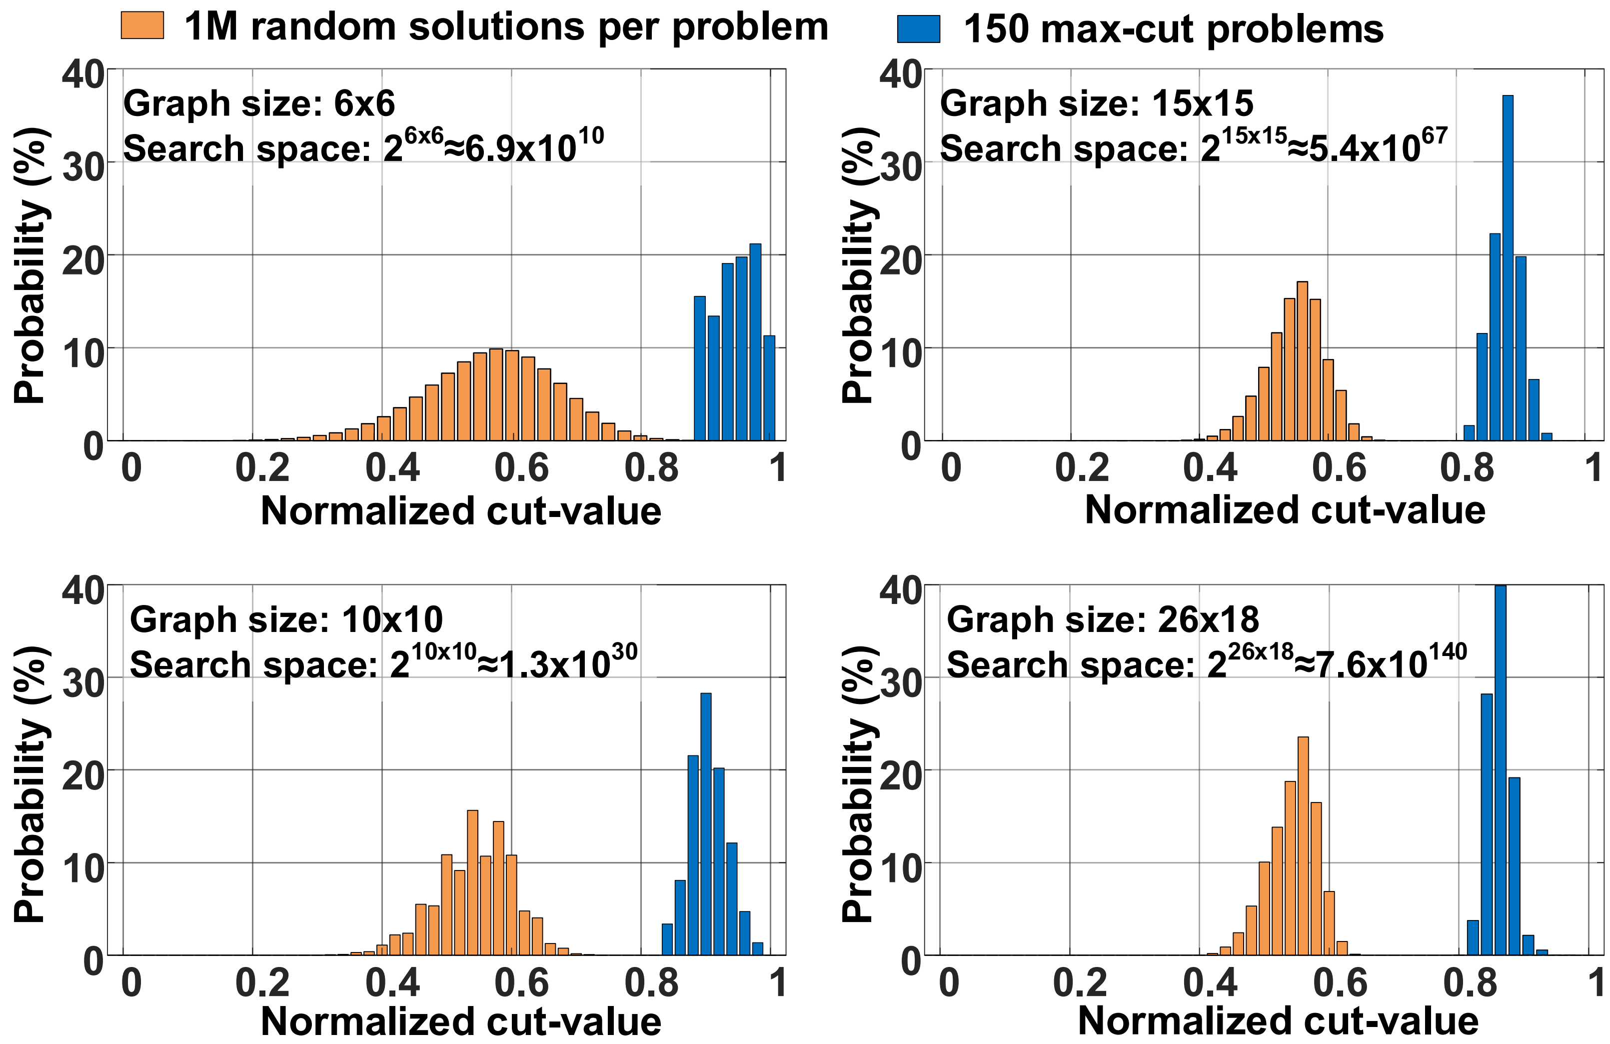

- Consistently better max-cut solution compared to 1 million random search

Comparison with Prior Arts

	PRL '19 [1]	IEDM '19 [2]	ISSCC '19 [3]	ISSCC '20 [4]	Quantum computer [5]	This work
Architecture	All to all	All to all	Near neighbor	Near neighbor	Near neighbor	Near neighbor
Technology	Photonics	IMT devices	CMOS 40nm	CMOS 65nm	Superconductor	CMOS 65nm
Coupling	Light modulation	IMT interaction	Digital logic	Digital logic	Qubit interaction	ROSC interaction
Operating temperature	Room temperature	Room temperature	Room temperature	Room temperature	-273.14°C	Room temperature
Total measured solution	16	1	10	1	Many	1000
Delay	1000 cycles	1ms	22μs	30 cycles	-	1μs - 10μs (est.)
Peak power	Not reported	Not reported	Not reported	Not reported	25KW	23mW
Measured accuracy	>95% (87% cases)	97.6% (4 spins)	98.8%	100% (Easy COP)	-	98%-100% (Easy) 82%-100% (Diff.)

[1] D. Pierangeli, PRL, 2019 [2] S. Dutta, IEDM, 2019 [3] T. Takemoto, ISSCC, 2019
[4] Y. Su, ISSCC, 2020 [5] Z. Bian, D-Wave Systems, 2010

Outline

- Motivation and Prior Art
- Circuit and Architecture Design
- Mapping Graph Problems to Hardware
- Measured Results
- Conclusion

Conclusion

- First hardware demonstration of a true coupling based integrated CMOS Ising computer
- Probabilistic exploration of various local minima
- Mapped and solved 1000 COPs in the chip with an accuracy of 82%-100%