





A Novel Oscillator Ising Machine Coupling Scheme for High-Quality Optimization

Shreesha Sreedhara^(✉)  and Jaijeet Roychowdhury 

EECS Department, University of California at Berkeley, Berkeley, USA
{shreesha, jr}@berkeley.edu

Abstract. Oscillator Ising Machines (OIMs) are networks of coupled nonlinear oscillators that solve the NP-hard Ising problem heuristically. Conventionally, the oscillators in an OIM are coupled using resistors. However, the phase-domain properties of such couplers are unsatisfactory; resistively-coupled OIMs do not realize the optimization performance predicted by simulations of idealized OIMs. This has been a major hurdle impeding the development of high quality analog OIMs on integrated circuits. In this paper, we present a novel coupling scheme, the *sampling coupler*, that addresses this issue theoretically and practically. Essentially, a sampling coupler injects a current that depends on the *phase difference* between interacting oscillators. We prove analytically that using sampling couplers leads to idealized OIMs, abstracting away the waveforms and innate phase sensitivities of the oscillators. We evaluate sampling-coupler OIMs (using simulation) on a practically-important digital wireless communication problem and show that the performance is near-optimal. Sampling couplers therefore open up a way to implement practically feasible, high-performance analog OIMs using virtually any oscillator.

1 Introduction

Ising machines are hardware solvers for the Ising problem, a general mathematical formulation involving an energy-like quantity (the Ising Hamiltonian, a quadratic function of binary problem variables called spins). They have been a focus of research in recent years (*e.g.*, [1, 5, 11, 24, 26, 27, 29]), on account of their ability to solve combinatorial optimization (CO) problems using novel analog mechanisms. Virtually all CO problems can be mapped into Ising form [16], making them amenable to solution using Ising machines, which offer the promise of speed, energy efficiency and miniaturisability.

Oscillator Ising Machines (OIMs), a scheme based on the dynamics of suitably-designed networks of coupled oscillators, have shown particular promise for high-quality optimization [27, 28]. Achieving this promise in hardware is currently an important research thrust. So far, OIMs with idealized phase-domain functions ($F_c(\cdot)$, described below) have shown excellent optimization characteristics, in both simulation [23, 27, 28] and a custom digital IC emulator [22].

However, *integrated circuit realizations of OIMs that employ real analog oscillators to deliver high-quality optimizations on practically-important problems* have not been demonstrated yet. Beyond the intrinsic scientific value of such a demonstration, high-quality analog OIM implementations can offer important practical advantages over simulation and digital emulation, such as significantly lower energy-to-solution.

In this paper, we identify a fundamental difficulty in designing high-quality analog OIMs using prior schemes, and present a new OIM scheme that overcomes it. Existing OIM schemes rely on resistive couplings.¹ An OIM's overall mathematical model (the generalized Kuramoto model, explained below) is obtained by combining the resistive coupling equation with the oscillators' nonlinear phase-sensitivity function (the PPV, see below²). This combination results in the generalized-Kuramoto function $F_c(\cdot)$, the precise nature of which crucially determines how well the OIM solves optimization problems. A key problem with resistive coupling is that the resulting $F_c(\cdot)$ function essentially mirrors the oscillator's PPV function—which makes it very difficult, indeed practically impossible, to tailor its shape to improve the OIM's optimization performance. As a result, practical analog OIM hardware designs often do not match results from simulations/emulations, which use idealized $F_c(\cdot)$ functions that result in high-quality optimization performance.

Although it is difficult/impossible to alter oscillator designs to yield a desired $F_c(\cdot)$ shape with resistive coupling, an alternative possibility is to change the coupling mechanism to reach the same goal. We show here that this is not only possible, but easy to achieve in practice using simple circuitry. The new scheme is highly effective in enabling excellent OIM optimization performance using standard analog oscillator designs which do not perform well with resistive coupling.

The essential difference from resistive coupling is that the coupling signal injected into a target oscillator by a source oscillator depends not only on the source oscillator's waveform, but also on the target oscillator's waveform, in a multiplicative manner. In other words, multiplicative feedback from the target oscillator modifies the injection from the source oscillator. Because multiplication captures phase differences, the effect in the phase domain is an *injection that is dependent on the phase difference between the target and source oscillators*—in contrast to resistive coupling, where the injection depends only on the source oscillator's phase.

We present a simple circuit, which we dub the *sampling coupler*, that implements this new coupling scheme using early-late sampling. We prove that the new coupling scheme translates mathematically to a generalized-Kuramoto $F_c(\cdot)$ function that matches idealized ones that achieve excellent optimization quality

¹ More precisely: independent of the actual hardware implementation (which can use, e.g., active elements instead of resistors), an oscillator couples to another by injecting a signal proportional to its voltage waveform.

² PPV stands for Perturbation Projection Vector; it is a periodic function that completely captures the dynamics of an oscillator's phase response to external inputs, such as those from other oscillators via coupling [7].

in simulation/emulation. We then simulate the generalized-Kuramoto model for the new coupling scheme and show that it performs essentially-perfect optimization, in contrast to resistive coupling, which yields poor results. The problem set we use for validation is a large set of benchmarks for the MU-MIMO³ detection problem, an important practical problem that arises in modern digital wireless communications [20].

Our invention and validation of the sampling coupler constitutes a key advance for enabling genuinely analog IC realizations of high-performance OIMs in the short term. The concept may also have implications for other types of Ising machines.

The remainder of the paper is organized as follows. We provide background on OIMs and their mathematical models in Sect. 2. This is followed by development of the sampling coupler (Sect. 3). It constitutes the core of this paper. In Sect. 4, we present simulation results evaluating a sampling-coupler based OIM on MU-MIMO detection problems.

2 Background: OIM Mathematical Models

The Ising model is simply a weighted graph, *i.e.*, a collection of nodes/vertices and branches/edges between some pairs of nodes, with each branch having a real-number weight. Each node (termed a “spin” in this context) is allowed to take two values, either 1 or -1 . Associated with this graph is an expression, the **Ising Hamiltonian**, which multiplies the weight of each branch by the values of the two spins it connects to, and sums over all branches, *i.e.*,

$$H = -\frac{1}{2} \sum_{\substack{i,j \\ i \neq j}} J_{ij} s_i s_j, \quad \text{where } s_i \in \{-1, +1\} \quad (1)$$

are the spins. The Ising Hamiltonian is sometimes interpreted as an “energy” associated with a given configuration of the spins, although in many situations they have no connection with energy in physics. The “Ising problem” is to find spin configurations with the minimum possible energy.

In 2016, we proposed using networks of coupled oscillators to solve Ising problems [26]. In our scheme, each of the N spins of an Ising problem is implemented by an oscillator. The information needed to find a solution of the Ising problem is encoded in the phase of each oscillator. This purely classical scheme had a significant advantage over prior Ising machines: OIMs can potentially be implemented entirely on chip in CMOS device technologies, with all the attendant benefits of IC integration—small physical size, low power consumption, easy scalability to many spins and easy mass production at low cost.

A simple model of the phase dynamics of the oscillator network [7, 13, 26] is the Kuramoto model, which takes the form

$$\frac{1}{f_0} \frac{d}{dt} \Delta\phi_i(t) = -K_c \sum_{\substack{j=1 \\ j \neq i}}^N J_{ij} \sin(2\pi(\Delta\phi_i(t) - \Delta\phi_j(t))). \quad (2)$$

³ MU-MIMO stands for Multi-User Multi-Input Multi-Output.

Here, $\Delta\phi_i(t)$ is the phase change of the i^{th} oscillator due to the influences of the other oscillators via coupling, K_c is a positive constant, and f_0 is the nominal oscillation frequency, assumed the same for all the oscillators. Note that (2) is repeated for each oscillator, resulting in a coupled system of nonlinear differential equations describing the dynamics of the network.

This system is (as far as is known) impossible to solve analytically. However, a key result that underpins OIMs has been established analytically [26]. The result is that (2) can be equipped with a Lyapunov function [17], given by

$$E(\{\Delta\phi_k\}) \triangleq -K_c \sum_{i=1}^N \sum_{\substack{j=1 \\ j \neq i}}^N J_{ij} \cos(2\pi(\Delta\phi_i - \Delta\phi_j)). \quad (3)$$

The utility of a Lyapunov function is that it is non-increasing (*i.e.*, it always decreases, or remains constant) with time if the phases $\Delta\phi_i(t)$ obey (2). In other words, the Lyapunov function always decreases (eventually settling to a constant) as the oscillator system’s dynamics evolve. The minima of the Lyapunov function correspond to stable equilibrium points (DC solutions) of (2), *i.e.*, **the coupled oscillator system’s phases settle to values that minimize (locally) the Lyapunov function.**

The Lyapunov function (3) looks very similar to the Ising Hamiltonian (1) if the coupling weights are the same as (or a scaled version of) the Ising problem’s J_{ij} . The main difference is that the Ising Hamiltonian contains $s_i s_j$ terms, while the Lyapunov function has a $\cos(2\pi(\Delta\phi_i - \Delta\phi_j))$ term. But if $2\pi\Delta\phi_i$ and $2\pi\Delta\phi_j$ were **restricted to either 0 or π** , with 0 defined as a spin value of 1 and π of -1 , $\cos(2\pi(\Delta\phi_i - \Delta\phi_j))$ would equal $s_i s_j$, and the oscillator network’s Lyapunov function would simply become a scaled version of the Ising Hamiltonian. In other words, the oscillator network’s dynamics would innately solve the Ising problem, at least to the extent that it would find a local minimum of the Lyapunov function, which is a continuized version of the Ising Hamiltonian – if each oscillator’s phase could somehow be restricted to be either 0 or π .

In general, there is no guarantee that the oscillator’s phases will settle to either 0 or π —indeed, phases settle to steady-state values that range continuously over $[0, 2\pi]$. In such cases, the oscillator network’s dynamics do not lead it to solutions that correspond to minima of the Ising Hamiltonian. It is in this context that “binarizing” the oscillator’s phases using a phenomenon called SHIL⁴ becomes important, since it restricts each oscillator’s phase to 0 or π . Modifying the oscillator network with SHIL injection changes the Kuramoto equations to

⁴ SHIL is an abbreviation of Sub-Harmonic Injection Locking; it is a phenomenon observed in nonlinear oscillators. In SHIL, the oscillator gets forced to oscillate in either one of two stable phases separated by π radians when it is perturbed by an external signal of a frequency that is twice the natural frequency of the oscillator [2].

$$\frac{1}{f_0} \frac{d}{dt} \Delta\phi_i(t) = -K_c \sum_{\substack{j=1 \\ j \neq i}}^N J_{ij} \sin(2\pi(\Delta\phi_i(t) - \Delta\phi_j(t))) - K_s \sin(4\pi\Delta\phi_i(t)). \quad (4)$$

The corresponding Lyapunov function becomes

$$E(\{\Delta\phi_k\}) \triangleq -K_c \sum_{i=1}^N \sum_{\substack{j=1 \\ j \neq i}}^N J_{ij} \cos(2\pi(\Delta\phi_i - \Delta\phi_j)) - K_s \sum_{i=1}^N \cos(4\pi\Delta\phi_i), \quad (5)$$

i.e., an additional term $\frac{K_s}{2} \sum \cos(4\pi\Delta\phi_i)$ appears, where K_s represents the strength of the input for inducing SHIL (called henceforth as SYNC). Since SHIL forces $2\pi\Delta\phi_i$ to either 0 or π (approximately), this term represents simply an addition of a constant offset to the Ising Hamiltonian, leaving the minima unchanged. Thus, **the coupled oscillator network with the addition of a 2nd harmonic SYNC input to induce SHIL naturally settles to phase solutions that locally minimize the Lyapunov function.**

This result is currently the best-known theoretical basis for oscillator Ising machines.

The above result only guarantees that the oscillator network will find **local** minima. To get the network out of local minima and guide it towards the global minimum, additional steps are needed.⁵ An effective way to get the network out of local minima is to relax or remove the SYNC signal (which binarizes each oscillator's phases through SHIL) and restore it again several times. Reducing SYNC allows the oscillators to drift away from $0/\pi$ to continuous values; as SYNC is ramped up again, the system tends to find its way to minima that are lower than previous ones. Adding a moderate amount of noise to the system helps with the Lyapunov minimization process.

The sinusoidal functions in the above equations, originally proposed by Kuramoto [15], do not suffice for practical oscillators, which require the more general form⁶

$$\frac{1}{f_0} \frac{d}{dt} \Delta\phi_i(t) = \sum_{j=1, j \neq i}^n J_{ij} F_c(\Delta\phi_i(t) - \Delta\phi_j(t)) + F_s(2\Delta\phi_i(t)). \quad (6)$$

The key difference from the basic Kuramoto equations (4) is that the $\sin(2\pi(\cdot))$ functions in (4) are replaced by $F_c(\cdot)$ and $F_s(\cdot)$ in (6). These functions can be extracted from, *e.g.*, the detailed circuit description of an oscillator provided to circuit simulators for low-level electronic simulation [2, 8]. Extracting $F_c(\cdot)$

⁵ This is characteristic of all Ising machines, as well as of optimization algorithms like simulated annealing.

⁶ There is a corresponding Lyapunov function [25] that generalizes (5).

and $F_s(\cdot)$ from low-level circuit differential equations involves first finding an abstraction known as the PPV phase-domain model [7,8] for the oscillator, *i.e.*,

$$\frac{1}{f_0} \frac{d}{dt} \Delta\phi(t) = \mathbf{v}^T(f_0 t + \Delta\phi(t)) \cdot \mathbf{b}(t). \quad (7)$$

In (7), the quantity $\mathbf{v}(t)$, called the PPV, represents a “nonlinear sensitivity” of the oscillator’s phase response to input perturbations; it can be extracted from the detailed differential equations of any oscillator using numerical techniques [8]. Furthermore, $\mathbf{b}(t)$ represents ‘small’ perturbations (inputs) to the oscillator, in response to which the oscillator’s phase changes by $\Delta\phi(t)$ —obtained by solving (7). Note that the waveform of the perturbed oscillator (denoted by $\mathbf{x}(t)$) is given by

$$\mathbf{x}(t) = \mathbf{x}_s(f_0 t + \Delta\phi(t)), \quad (8)$$

where $\mathbf{x}_s(\cdot)$ represents the 1-periodic steady-state waveform of the unperturbed oscillator. $\Delta\phi(t)$ thus represents the additional phase shift due to the perturbation.

Once (7) is available, an averaging or “Adlerization” process [2,4] is used to extract $F_c(\cdot)$ and $F_s(\cdot)$. The derivation assumes that couplings between oscillators are resistive—*i.e.*, the injected signal into a target oscillator due to coupling with other oscillators is proportional to the waveform of a source oscillator. Using $F_c(\cdot)$ and $F_s(\cdot)$ in (6), OIM systems with many spins can be simulated quickly to assess Lyapunov/Hamiltonian minimization performance.⁷ We refer the reader to, *e.g.*, [27,28] as a starting point for further information about OIMs and their underlying mathematics.

3 Sampling Coupler

The precise nature of the 2π -periodic, typically non-sinusoidal, functions $F_c(\cdot)$ and $F_s(\cdot)$ strongly influences Hamiltonian minimization performance of the OIM. Unfortunately, function shapes that typically emerge from practical oscillator designs do not lead to high-quality minimization performance. Moreover, it is very difficult (indeed essentially impossible) to alter an oscillator design to achieve desired shapes for $F_c(\cdot)$ or $F_s(\cdot)$. In this work, we circumvent this difficulty by changing the coupling mechanism, and re-deriving (7) with this change to achieve $F_c(\cdot)$ and $F_s(\cdot)$ functions that do produce high-quality optimization performance in OIMs.

The new coupling scheme and corresponding circuit, termed the *sampling coupler*, results in a near-ideal square-wave shape for $F_c(\cdot)$ or $F_s(\cdot)$.⁸ The square-wave shape for $F_c(\cdot)$ is desirable because, empirically, we have observed that

⁷ Simulating a coupled system of “Un-Adlerized” PPV equations for the OIM network, as we do to generate some of our results in this paper, provides more accurate results than (6), though it requires somewhat greater computational effort.

⁸ For simplicity and brevity, we focus on $F_c(\cdot)$ in the following; the reasoning for $F_s(\cdot)$ is very similar.

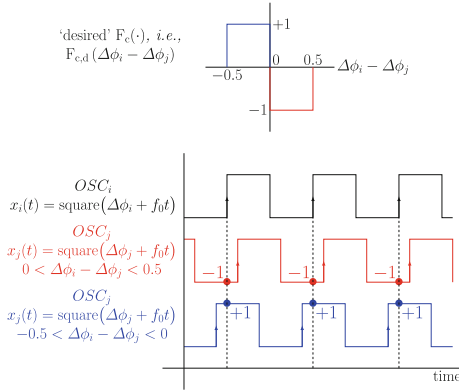


Fig. 1. (Top) A plot of ideal ‘sharp’, *i.e.*, the desired $F_c(\Delta\phi_i - \Delta\phi_j)$ denoted by $F_{c,d}(\Delta\phi_i - \Delta\phi_j)$. (Bottom) Example waveforms of OSC_i (black) and OSC_j (blue and red). The waveform in red is for the case when $0 < \Delta\phi_i - \Delta\phi_j < 0.5$, and the waveform in blue is for the case when $-0.5 < \Delta\phi_i - \Delta\phi_j < 0$. OSC_j ’s waveform (red and blue) are sampled at the rising edges of OSC_i ’s waveform (black). The samples (shown as bullets) directly provide the values of $F_{c,d}(\Delta\phi_i - \Delta\phi_j)$.

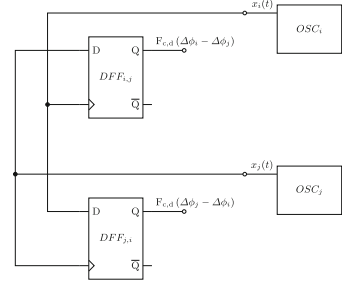


Fig. 2. Flip-flops to directly evaluate $F_{c,d}(\cdot)$. $DFF_{i,j}$ samples the waveform of OSC_j at the rising edges of OSC_i , which gives us $F_{c,d}(\Delta\phi_i - \Delta\phi_j)$. Similarly, samples of $DFF_{j,i}$ provide $F_{c,d}(\Delta\phi_j - \Delta\phi_i)$.

it leads to very good optimization performance (*e.g.*, see Fig. 6); other shapes, *e.g.*, those that emerge naturally for ring and other oscillators, lead to significant performance degradation. Another significant practical advantage of a square-wave $F_c(\cdot)$ is that simple circuit structures suffice for implementation.

Figure 1 introduces the concept and explains the choice of the term “sampling coupler”. The upper figure depicts the above-mentioned desired square-wave shape of $F_c(\cdot)$, which we refer to as $F_{c,d}(\cdot)$ henceforth,⁹ to distinguish it from the actually-achieved $F_c(\cdot)$ for a given OIM scheme. Because of this shape, note that any term $F_{c,d}(\Delta\phi_i - \Delta\phi_j)$ in (6) takes only 2 values, ± 1 . If $\Delta\phi_i - \Delta\phi_j$ is restricted to a single period $[-0.5, 0.5]$, $F_{c,d}(\Delta\phi_i - \Delta\phi_j)$ is -1 if $\Delta\phi_i > \Delta\phi_j$, and $+1$ if $\Delta\phi_i < \Delta\phi_j$. In other words, the value of the term depends only on whether *one phase is ahead of, or behind, the other*.

Recall that $\Delta\phi_i$ and $\Delta\phi_j$ are the phases of oscillatory waveforms. If these are square as well,¹⁰ as shown in the lower part of Fig. 1, then *simply looking at (or sampling) the value of one waveform at the transition edge of the other* suffices to determine if the phase of one is ahead of, or behind, the other. This

⁹ The ‘d’ stands for ‘desired’.

¹⁰ It is easy in practice to turn most waveforms into square ones using a simple thresholding circuit.

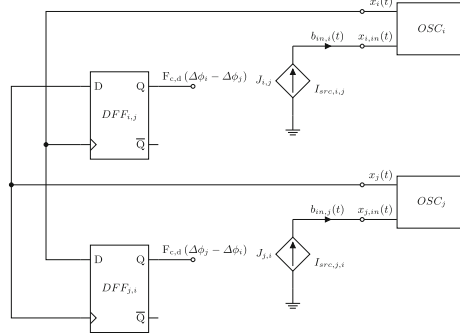


Fig. 3. A two oscillator OIM with sampling couplers. OSC_i couples to OSC_j as follows: the lower flip-flop samples the square-wave output of OSC_i at the transition edge of OSC_j ; the resulting sampled value of ± 1 is weighted by the Ising coupling weight $J_{j,i}$ and injected into OSC_j . The upper flip-flop couples OSC_j to OSC_i in the same manner. Note that $J_{j,i} = J_{i,j}$ for the Ising model to be valid.

is called *early-late sampling*; we have just established that it achieves evaluating $F_{c,d}(\Delta\phi_i - \Delta\phi_j)$. A block-level circuit implementing this using D flip-flops is shown in Fig. 2; the outputs Q produce $F_{c,d}(\Delta\phi_i - \Delta\phi_j)$.

Now consider the following scheme. Ignoring the provenance of the circuit of Fig. 2 as explained above, suppose we use its outputs Q to couple oscillators, as shown in Fig. 3. Coupling from a source oscillator j to a target oscillator i is effected by sampling the source oscillator's square-wave signal at the transition of the target oscillator's square-wave output; then, weighting the ± 1 sampled value (which is equal to $F_{c,d}(\Delta\phi_i - \Delta\phi_j)$) by the Ising coupling weight $J_{j,i}$; and finally, injecting a current equal to this value into the target oscillator. Generalization to a system of N coupled oscillators (Fig. 4) is straightforward.

It now remains to understand what an OIM with such a coupling scheme can achieve. Indeed, is the generalized Kuramoto model (6) still applicable to this scheme? If so, what does $F_c(\cdot)$ in (6) turn out to be?

Our key result is that (6) is still valid, with an $F_c(\cdot)$ function that is a scaled version of the desired $F_c(\cdot)$ illustrated in Fig. 1, *i.e.*, $F_c(\cdot) = K_c \cdot F_{c,d}(\cdot)$ for some constant K_c . Crucially, the square-wave shape of the achieved $F_c(\cdot)$ essentially *does not depend on the oscillator's PPV $v(\cdot)$ in (7)*, except for the scaling factor K_c . This feature is in stark contrast to the resistive/proportional coupling scenario, where the shape of $F_c(\cdot)$ is determined to a great extent by that of $v(\cdot)$ —leading to sub-optimal Hamiltonian minimization performance. The derivation/proof of these results, essentially a specialized version of the Adlerization/averaging procedure of [2], is provided in Appendix A.

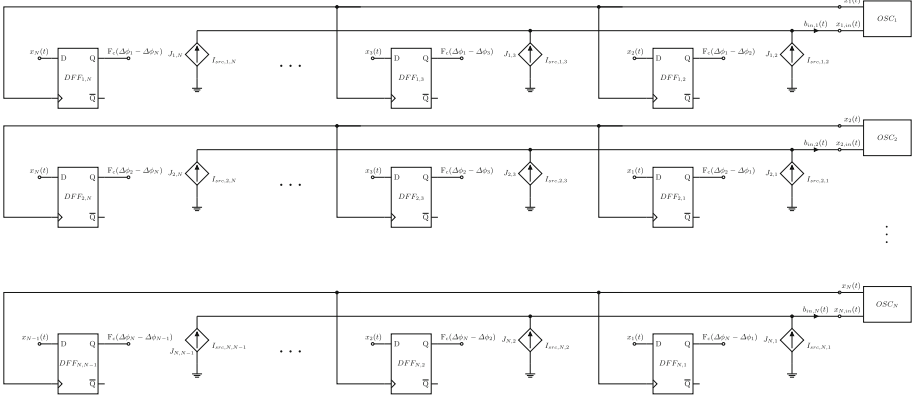


Fig. 4. An N oscillator OIM that employs sampling couplers. Note that (for every i) a sampling coupler that couples the i^{th} oscillator to itself does not exist.

4 Results

We now present results comparing the optimization performance of OIM with resistive coupling with that of OIM using the new sampling coupler scheme. Result data for OIM with the sampling coupler, and with resistive coupling, was generated by simulating a system of PPV equations (7) incorporating the appropriate coupling mechanism. Using the PPV equations for simulating the sampling-coupler and resistor based OIMs is closer in terms of fidelity to actual circuit implementations compared to generalized Kuramoto models. For comparison, the generalized Kuramoto model (6) using the ideal square-wave desired $F_{c,d}(\cdot)$ was also simulated.¹¹

A large set of benchmarks (550,000 problems, [20, 22, 23]) for MU-MIMO detection [10, 14], an important practical problem in wireless telecommunications, was used to evaluate performance. We first describe the MU-MIMO detection problem before presenting our simulation results.

Modern wireless communication settings involve multiple users with single/multiple transmit antennæ, using the same resources (time and frequency)




Fig. 5. An illustration of a Multiple User Multiple In Multiple Out (MU-MIMO) setup. Multiple users transmit their data, \mathbf{x} , to a receiver with multiple antennæ, where the signal \mathbf{y} is measured. Transmission occurs over several paths, characterized by the channel matrix H .

¹¹ The simulations used the Forward Euler technique for numerical solution of ordinary differential equations, as noted in [26].

to transmit to a receiver equipped with multiple receive antennæ (see Fig. 5).¹² As a result, each received signal consists of a noisy superposition of several users’ transmitted symbols; the influence on each receiving antenna of each transmitting antenna is captured mathematically using a *channel matrix*, H . Recovering the originally-sent symbols from received signals involves solving a hard combinatorial optimization problem (the MU-MIMO detection problem [6, 14]) to infer the most likely set of transmitted symbols, given the set of signals received. Solving exactly for the most likely transmitted symbols, *i.e.*, the M.L. (Maximum Likelihood) solution, is too computationally expensive to be practical. In practical 4G/5G/6G systems, it is necessary to meet real-time performance requirements during detection. Therefore, heuristic methods that use much less computation than M.L., such as LMMSE (“linear minimum mean-square error”) or ZF (“Zero Forcing”) [19], are universally employed even though they do not recover transmitted symbols as accurately as M.L.. However, even though impractically expensive in practice, results from M.L. serve as the gold standard for performance, *i.e.*, they represent the best SER (symbol error rate) achievable, via optimal detection.

The MU-MIMO detection problem can be cast as an Ising (Hamiltonian minimization) problem, making it possible for Ising machines, such as OIMs, to solve it. If the number of transmit antennæ is denoted by N_t , it can be shown ([20, Section 4], [14]) that the equivalent Ising problem in the form (1) requires $N_t + 1$ spins. The coupling coefficients J_{ij} are typically all non-zero, *i.e.*, every spin is non-trivially coupled to every other spin; such all-to-all coupling is termed *dense*. Finding the global minimum of the Ising version of the problem corresponds exactly to finding the set of transmit symbols most likely to result in a given observed set of received signals, *i.e.*, the M.L. solution. A sub-optimal solution is typically characterized by a larger number of symbol errors than in the optimal solution. The fidelity of any detection scheme can be judged by the SER it achieves, compared to the optimal SER achieved by M.L..

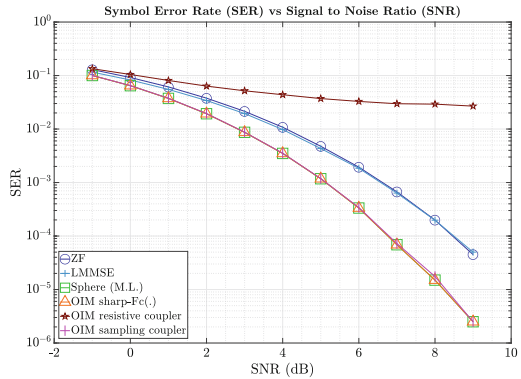


Fig. 6. SER vs SNR plot of a resistive coupler OIM and a sampling coupler OIM compared against the idealized OIM and other heuristic detection schemes. As can be seen, the performance of the resistive coupler OIM is very poor when compared to the optimal decoder (Sphere). The sampling coupler OIM matches the SER of the Sphere decoder.

¹² Doing so greatly increases data rate and reliability compared to single-antenna systems.

Figure 6 depicts SER results from detection using several different schemes: M.L. (using an efficient implementation known as Sphere Decoding [9, 12, 18]), LMMSE and ZF, OIM with resistive coupling, and OIM with our new sampling coupler. The benchmark set used for our simulations consists of 11 sets of test problems. Each set corresponds to a specific SNR (signal-to-noise ratio) at the receiving antennæ; the 11 sets of problems have SNR (in dB, *i.e.*, $10 \log_{10}(\text{actual SNR})$) varying from -1 dB to 9 dB, shown on the horizontal axis of Fig. 6. For each SNR value, the test problem set consists of 1000 different channel matrices H , for each of which 50 pairs of transmitted symbol vectors \mathbf{x} and received signal vectors \mathbf{y} are available. Thus, there are 50,000 decoding problems for each SNR value, or a total of 550,000 problems in all. The vertical axis represents (also on a logarithmic scale) the SERs achieved by each detection scheme across all 50,000 problems for each value of SNR. As can be seen, problem sets with higher SNRs generally have fewer symbol errors; those with lower SNRs have more. This is intuitive because more noise leads (probabilistically) to more bit/symbol upsets.

The performances of the different detection techniques are clustered into 3 groups, as can be seen. The lowest cluster comprises 3 traces: the Sphere Decoder (representing optimal SER performance), OIM using the sampling coupler (simulated using the more detailed PPV equations (7), as noted earlier), and OIM simulated at the generalized Kuramoto level (6) using the ideal square-wave $F_{c,d}(\cdot)$ —this was the starting point of our development of the sampling coupler in Sect. 3. The SER performance of both OIM simulations are essentially identical to the optimal one from the Sphere Decoder/M.L., *i.e.*, sampling-coupler-based OIM delivers the best possible SER performance on these problems.

The two traces in the middle show the SER performances of LMMSE and ZF. It is apparent that they are significantly inferior to the three methods above. The uppermost trace shows the SER performance of a typical OIM with resistive coupling, also simulated at the more-accurate PPV level. Its SER performance is inferior even with respect to LMMSE and ZF, and far inferior to that of OIM with the sampling coupler. These results underscore the potential impact of using the sampling coupler for practical OIM designs.¹³

5 Conclusion

Virtually all current OIM schemes use resistors to couple oscillators. However, due to their strong dependence on oscillator waveforms and PPVs, their characteristic coupling functions $F_c(\cdot)$ tend to be smooth and skewed. These shapes impact the performance of resistively-coupled OIMs, which often is significantly inferior to that of OIMs with idealized square-wave $F_c(\cdot)$ s. In this paper, we have presented a novel coupling scheme, the sampling coupler, which directly injects a current that is a function of the idealized square-wave $F_{c,d}(\cdot)$. Moreover, we have shown a simple flip-flop based circuit structure that implements

¹³ Note, however, that resistive coupling in OIMs with specially designed oscillators can produce good results for some problems, such as MAX-CUT [21].

the sampling coupler. We have proved analytically that the actual $F_c(\cdot)$ achieved by this coupling circuit/scheme is equal to a scaled version of the desired ideal square-wave $F_{c,d}(\cdot)$; the impact of the oscillators' waveforms and the PPVs is reduced to a mere scaling factor. Furthermore, using simulations, we have shown that sampling-coupler based OIMs perform near-optimally on practically relevant MU-MIMO detection problems. Sampling couplers thus constitute an important technology for realizing practically feasible high-performance analog OIMs that are insensitive to details of oscillator implementation.

Acknowledgments. This work was enabled by support from the US Defense Advanced Research Projects Agency (DARPA) and the US National Science Foundation (NSF); additional support was provided by Berkeley's Bakar Prize Program. The MU-MIMO benchmark set used in this paper was provided by Pavan K. Srinath and Joachim Wabnig of Nokia Bell Labs. We thank Thomas Hart and Charles Macedo for important inputs.

A Sampling-Coupler Based OIMs Achieve $F_{c,d}(\cdot)$

Here, we prove that the actual $F_c(\cdot)$ of the sampling coupler is in fact equal to a scaled version of the ideal $F_{c,d}(\cdot)$. We start from the PPV equation ((7), repeated here):

$$\frac{1}{f_0} \frac{d}{dt} \Delta\phi(t) = \mathbf{v}^T(f_0 t + \Delta\phi(t)) \cdot \mathbf{b}(t). \quad (9)$$

Here, $\Delta\phi(t)$ is the phase change of the oscillator, $\mathbf{v}^T(\cdot)$ is its vector of 1-periodic PPVs, and $\mathbf{b}(t)$ is the vector of inputs applied to the oscillator.

Applying the above PPV model to the two oscillators coupled via sampling couplers, we get

$$\begin{aligned} \frac{1}{f_0} \frac{d}{dt} \Delta\phi_i(t) &= v^T(f_0 t + \Delta\phi_i(t)) \cdot b_i(t), \\ \frac{1}{f_0} \frac{d}{dt} \Delta\phi_j(t) &= v^T(f_0 t + \Delta\phi_j(t)) \cdot b_j(t), \end{aligned} \quad (10)$$

where $b_i(t)$ and $b_j(t)$ represent the inputs into each oscillator from the other via the sampling couplers in Fig. 3. Note that vector PPVs and inputs (*i.e.*, $\mathbf{v}(\cdot)$ and $\mathbf{b}(t)$ respectively) have been replaced by scalars; this is a simplification (for exposition) assuming a single scalar input, *i.e.*, \mathbf{b} has only one nonzero component.¹⁴

¹⁴ In (10), $v(\cdot)$ and $b_i(t)$ are the PPV and the input (respectively) of the node $x_{i,in}$ in Fig. 3.

We now focus on OSC_i and derive its actual $F_c(\cdot)$. The input $b_i(t)$, represented in Fig. 3 by $I_{src,i,j}$, has the form

$$b_i(t) = J_{i,j} \cdot F_{c,d}(\Delta\phi_i(t) - \Delta\phi_j(t)) \cdot w(\Delta\phi_i(t) + f_0t), \quad (11)$$

where:

- $J_{i,j}$ is the Ising coupling coefficient (from (1)) between the i^{th} and the j^{th} oscillator.
- $F_{c,d}(\Delta\phi_i(t) - \Delta\phi_j(t))$ is the value of the sample held by $DFF_{i,j}$ in Fig. 3, as established in Sect. 3. Note that $\Delta\phi_i$ and $\Delta\phi_j$ in (11) now change with time as the system evolves. However, the flip-flops in Fig. 3 hold the value at the last sampling instant until the next sample; this is not captured by $F_{c,d}(\Delta\phi_i(t) - \Delta\phi_j(t))$.
- The $w(\cdot)$ term captures the sampling aspect of the flip-flop. Note that the flip-flop $DFF_{i,j}$ samples at transitions of $x_i(t)$, *i.e.*, its sampling instant is timed using the phase $f_0t + \Delta\phi_i(t)$ of OSC_i . Ideal sampling would be captured by a weighted delta function of this phase, *i.e.*, $Cw(f_0t + \Delta\phi_i(t) + \theta)$, where $w(\cdot)$ is a unit impulse train with period 1, C is a weight, and θ is a constant phase offset, useful for adjusting the sampling instant within each cycle and/or to model clock-to-Q delay in the flip-flop.¹⁵ However, $w(\cdot)$ can in fact be almost any 1-periodic function for the scheme to work, as we show below; incorporating C into, and using a θ -shifted version of, a given $w(\cdot)$ simplifies the expression to $w(f_0t + \Delta\phi_i(t))$.

Substituting $b_i(t)$ in (10), we obtain

$$\begin{aligned} \Delta\dot{\phi}_i(t) = & f_0 \cdot v(\Delta\phi_i(t) + f_0t) \cdot J_{i,j} \\ & \cdot F_{c,d}(\Delta\phi_i(t) - \Delta\phi_j(t)) \cdot w(\Delta\phi_i(t) + f_0t). \end{aligned} \quad (12)$$

Now, we assume that the phases $\Delta\phi_i(t)$ and $\Delta\phi_j(t)$ vary ‘slowly’—this is a standard assumption for averaging or “Adlerization”; [2,4]. With this assumption, the Adlerization of (12) is

$$\begin{aligned} \Delta\dot{\phi}_i(t) \approx & f_0 \int_{\phi=0}^{\phi=1} v(\Delta\phi_i(t) + \phi) \cdot J_{i,j} \\ & \cdot F_{c,d}(\Delta\phi_i(t) - \Delta\phi_j(t)) \cdot w(\Delta\phi_i(t) + \phi) \cdot d\phi, \end{aligned} \quad (13)$$

¹⁵ The flip-flop can have a nonzero clock-to-Q delay, but we assume it is designed to avoid metastability, *i.e.*, such delays will not be indefinitely long.

where ϕ represents the nominal oscillator phase, $f_0 t$. Simplifying the above leads to

$$\begin{aligned}
 \Delta \dot{\phi}_i(t) &= f_0 \cdot J_{i,j} \cdot F_{c,d}(\Delta \phi_i(t) - \Delta \phi_j(t)) \cdot \\
 &\quad \int_{\phi=0}^{\phi=1} v(\Delta \phi_i(t) + \phi) \cdot w(\Delta \phi_i(t) + \phi) \cdot d\phi \\
 &= f_0 \cdot J_{i,j} \cdot F_{c,d}(\Delta \phi_i(t) - \Delta \phi_j(t)) \int_{\psi=\Delta \phi_i(t)}^{\psi=1+\Delta \phi_i(t)} v(\psi) \cdot w(\psi) \cdot d\psi \tag{14} \\
 &= f_0 \cdot J_{i,j} \cdot F_{c,d}(\Delta \phi_i(t) - \Delta \phi_j(t)) \int_{\psi=0}^{\psi=1} v(\psi) \cdot w(\psi) \cdot d\psi,
 \end{aligned}$$

where $\psi \triangleq \Delta \phi_i(t) + \phi$. For the last step, we used the fact that the integral remains the same over any interval of length 1 since both $v(\cdot)$ and $w(\cdot)$ are 1-periodic.

It is convenient, though not necessary,¹⁶ to assume that

$$K_c \triangleq \int_{\phi=0}^{\phi=1} v(\phi) \cdot w(\phi) \cdot d\phi > 0. \tag{15}$$

Using this definition of K_c in (14), we get

$$\Delta \dot{\phi}_i(t) = f_0 \cdot J_{i,j} \cdot K_c F_{c,d}(\Delta \phi_i(t) - \Delta \phi_j(t)). \tag{16}$$

Comparing (16) to (6), we have shown that the actual $F_c(\cdot)$ from using sampling couplers is indeed equal to a scaled version of $F_{c,d}(\theta)$.

The above generalizes straightforwardly to the case of N coupled oscillators, resulting in

$$\forall i, \quad \Delta \dot{\phi}_i(t) = f_0 \sum_{j=1, j \neq i}^{j=N} J_{i,j} \cdot K_c F_{c,d}(\Delta \phi_i(t) - \Delta \phi_j(t)). \tag{17}$$

¹⁶ If K_c is negative, then the signs of the coupling coefficients get reversed. The improbable case of K_c being zero or very small can be remedied by shifting $w(\cdot)$ by some delay.

References

1. Marandi, A., Wang, Z., Takata, K., Byer, R.L., Yamamoto, Y.: Network of time-multiplexed optical parametric oscillators as a coherent Ising machine. *Nat. Photonics* **8**(12), 937–942 (2014)
2. Neogy, A., Roychowdhury, J.: Analysis and design of sub-harmonically injection locked oscillators. In: *Proceedings of IEEE DATE*, March 2012
3. Adler, R.: A study of locking phenomena in oscillators. *Proc. I.R.E. Waves Elect.* **34**, 351–357 (1946)
4. Adler, R.: A study of locking phenomena in oscillators. *Proc. IEEE* **61**, 1380–1385 (1973), reprinted from [3]
5. Afoakwa, R., Zhang, Y., Vengalam, U.K.R., Ignjatovic, Z., Huang, M.: BRIM: bistable resistively-coupled Ising machine. In: *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 749–760. IEEE (2021)
6. Damen, M.O., El Gamal, H., Caire, G.: On maximum-likelihood detection and the search for the closest lattice point. *IEEE Trans. Inf. Theory* **49**(10), 2389–2402 (2003)
7. Demir, A., Mehrotra, A., Roychowdhury, J.: Phase noise in oscillators: a unifying theory and numerical methods for characterization. *IEEE Trans. Ckts. Syst. I.* **47**(5), 655–674 (2000). <https://doi.org/10.1109/81.847872>
8. Demir, A., Roychowdhury, J.: A reliable and efficient procedure for oscillator PPV computation, with phase noise macromodelling applications. *IEEE Trans. CAD* **22**, 188–197 (2003)
9. Fincke, U., Pohst, M.: Improved methods for calculating vectors of short length in a lattice, including a complexity analysis. *Math. Comput.* **44**(170), 463–471 (1985). <https://doi.org/10.2307/2007966>
10. Goutay, M., Aoudia, F.A., Hoydis, J.: Deep Hypernetwork-based MIMO Detection. In: *21st International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pp. 1–5. IEEE (2020)
11. Hamerly, R., et al.: Experimental investigation of performance differences between coherent Ising machines and a quantum Annealer. *Sci. Adv.* **5**(5), eaau0823 (2019)
12. Hassibi, B., Vikalo, H.: On the expected complexity of sphere decoding. In: *Conference Record of Thirty-Fifth Asilomar Conf. on Signals, Systems and Computers*, vol. 2, pp. 1051–1055. IEEE (2001)
13. Acebrón, J.A., Bonilla, L.L., Vicente, C.J.P., Ritort, F., Spigler, R.: The Kuramoto model: a simple paradigm for synchronization phenomena. *Rev. Mod. Phys.* **77**(1), 137 (2005)
14. Kim, M., Venturelli, D., Jamieson, K.: Leveraging quantum annealing for large MIMO processing in centralized radio access networks. In: *Proceedings of the ACM Special Interest Group on Data Communication*, pp. 241–255. ACM (2019)
15. Kuramoto, Y.: Self-entrainment of a population of coupled non-linear oscillators. In: Araki, H. (eds.) *International Symposium on Mathematical Problems in Theoretical Physics*. LNP, vol. 39, pp. 420–422. Springer, Berlin, Heidelberg (1975). <https://doi.org/10.1007/BFb0013365>
16. Lucas, A.: Ising formulations of many NP problems. *Front. Phys.* **2**, 5 (2014)
17. Lyapunov, A.M.: The general problem of the stability of motion. *Int. J. Control* **55**(3), 531–534 (1992)
18. Pohst, M.: On the computation of lattice vectors of minimal length, successive minima and reduced bases with applications. *SIGSAM Bull.* **15**(1), 37–44 (1981). <https://doi.org/10.1145/1089242.1089247>

19. Proakis, J.G., Salehi, M.: Digital Communications, 5th edn. McGraw Hill, Boston (2008)
20. Roychowdhury, J., Wabnig, J., Srinath, K.P.: Performance of Oscillator Ising Machines on Realistic MU-MIMO Decoding Problems. Research Square preprint (Version 1), September 2021
21. Sreedhara, S., Roychowdhury, J.: Oscillatory Neural Networks—from Devices to Computing Architectures, chap. Oscillator Ising Machines. Springer (2024). (to appear)
22. Sreedhara, S., Roychowdhury, J., Wabnig, J., Srinath, K.P.: Digital emulation of oscillator Ising machines. In: Proceedings of IEEE DATE, pp. 1–2 (2023)
23. Sreedhara, S., Roychowdhury, J., Wabnig, J., Srinath, P.K.: MU-MIMO Detection Using Oscillator Ising Machines. In: Proc. ICCAD. pp. 1–9 (2023)
24. Inagaki, T., et al.: A coherent Ising machine for 2000-node optimization problems. *Science* **354**(6312), 603–606 (2016)
25. Wang, T., Roychowdhury, J.: OIM: Oscillator-based Ising Machines for Solving Combinatorial Optimisation Problems. [arXiv:1903.07163](https://arxiv.org/abs/1903.07163) (2019)
26. Wang, T., Roychowdhury, J.: Oscillator-based Ising Machine. [arXiv:1709.08102](https://arxiv.org/abs/1709.08102) (2017)
27. Wang, T., Roychowdhury, J.: OIM: oscillator-based Ising machines for solving combinatorial optimisation problems. In: McQuillan, I., Seki, S. (eds.) UCNC 2019. LNCS, vol. 11493, pp. 232–256. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-19311-9_19
28. Wang, T., Wu, L., Nobel, P., Roychowdhury, J.: Solving combinatorial optimisation problems using oscillator based Ising machines. *Nat. Comput.* **20**, 1–20 (2021)
29. Bian, Z., Chudak, F., Israel, R., Lackey, B., Macready, W.G., Roy, A.: Discrete optimization using quantum annealing on sparse Ising models. *Front. Phys.* **2**, 56 (2014)