

COCOON - A Conductive Substrate-based Coupled Oscillator Network for Wireless Communication

Xingda Chen, Deepak Ganesan, Jeremy Gummeson, Mohammad Rostami

University of Massachusetts Amherst, Amherst, MA

xingdachen@umass.edu, dganesan@cs.umass.edu, jgummeso@umass.edu, mrostami@cs.umass.edu

ABSTRACT

Advances in flexible conductive substrates such as conductive wall-paper and paint present new opportunities for optimizing the performance of IoT nodes in smart homes and buildings. In this paper, we explore an unconventional use of such substrates for pulling frequencies of oscillators across IoT devices and wireless front-ends connected to the substrate. We show that by using this technique, we can replace precise crystal oscillators by lower precision and lower cost ceramic oscillators without compromising their ability to be used for tasks that require precise frequencies such as frequency-synchronized multi-static backscatter and synchronized sampling. We present an end-to-end design including a) analysis of conditions under which frequency pulling of oscillators across conductive substrates can work, b) a new technique to detect frequency locking across oscillators without requiring explicit communication, and c) an adaptive method that can be used to synchronize oscillators at minimum power consumption. We then show that these elements can be composed to design a high-performance multi-static backscatter system that performs as well as one that uses a shared high-precision clock but at an order of magnitude less monetary cost. We show that our system can scale and operate at very low power, while having low complexity since it requires no explicit interaction among devices attached to the substrate.

CCS CONCEPTS

• **Networks** → **Network architectures; Wireless access points, base stations and infrastructure.**

KEYWORDS

Synchronization, Smart textile, Novel radio network

ACM Reference Format:

Xingda Chen, Deepak Ganesan, Jeremy Gummeson, Mohammad Rostami. 2021. COCOON - A Conductive Substrate-based Coupled Oscillator Network for Wireless Communication. In *The 19th ACM Conference on Embedded Networked Sensor Systems (SenSys'21)*, November 15–17, 2021, Coimbra, Portugal. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3485730.3485940>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

SenSys'21, November 15–17, 2021, Coimbra, Portugal

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-9097-2/21/11...\$15.00

<https://doi.org/10.1145/3485730.3485940>

1 INTRODUCTION

Flexible conductive substrates have seen significant interest in recent years due to their potential to be cheaply manufactured using additive or roll-to-roll manufacturing and their ability to conform to three dimensional surfaces. A wide range of technologies are being explored in this space including conductive ink [4] i.e. paints infused with conductive particles such as silver and carbon [4], conductive threads that are produced by applying conductive materials such as silver, copper, or conductive polymers on fibers and yarns [7], and others such as conductive paper.

While the use of such substrates has been extensively explored from an interactions perspective, there are many unexplored opportunities from a distributed sensor systems perspective. For example, prior work has explored the use of a conductive wall as a sensor [8, 46] or as a living surface that responds to user touch inputs [41, 42]. But many other opportunities remain unexplored. For example, these substrates can possibly be leveraged to design low-cost yet sophisticated distributed RF systems that can operate in a synchronized manner for multi-static backscattering, beam forming and RF power delivery. We may also be able to enable low-cost synchronized arrays of microphones or imagers that are interconnected with conductive substrates. More generally, we argue that the widespread availability of conductive substrates can allow us to re-think how distributed sensing and communication applications are deployed in smart homes and buildings.

Sharing a Virtual Oscillator: In this paper, we explore the use of conductive substrates to re-think a fundamental building block of such applications i.e. low-cost yet high-precision synchronization. To achieve this, we view conductive substrates as a carrier of low-frequency analog signals to couple oscillators on different sensor nodes. It is well-known that oscillator circuits can couple when connected to each other, as long as the coupling meets certain criteria [18]. This implies that by simply attaching oscillators on different nodes to a low-resistance conductive substrate, we can potentially allow them to couple with each other and pull their clocks to the same frequency.

The ability to couple oscillators allows a group of devices that are connected to the same substrate to *share a high-precision virtual oscillator* despite being equipped with a low-precision oscillator. Oscillators trade precision for cost (see Figure 1) — oscillators such as a Temperature-Controlled Crystal Oscillator (TCXO) are very expensive and power-hungry but offer high precision even under temperature variations whereas ceramic and ring oscillators have very high variability and drift but lower cost [5]. IoT devices have to compromise between these two metrics with quartz crystal oscillators that offer a middle ground in terms of precision and cost. Sharing a virtual oscillator over a conductive substrate offers the best of both worlds since we can consider topologies where one

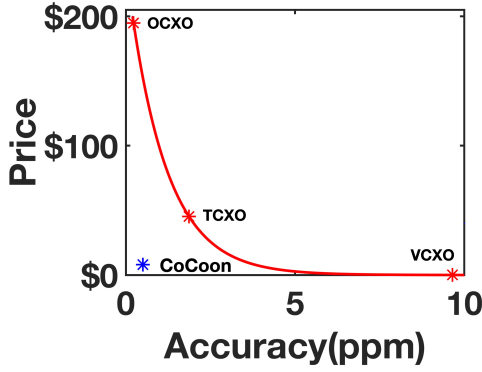


Figure 1: More accurate temperature controlled oscillators are expensive. COCOON enables cheap oscillators such as voltage-controlled crystal oscillators (VCXO) or ceramic oscillators to have high accuracy.

node on the substrate has a more precise and more expensive master oscillator such as a TCXO that synchronizes a number of poor quality and very low-cost but cheap slave oscillators (e.g. ceramic oscillators) by leveraging coupling.

A key advantage of this approach is lower cost and complexity of deploying sophisticated communication systems. For example, we can develop multi-static backscatter communication systems (e.g. a network of RFID readers) that require tight synchronization across transmitters and receivers but at a fraction of the cost since synchronization overhead is offloaded onto the substrate. Sharing a virtual high-precision oscillator can also greatly mitigate Carrier Frequency Offset (or CFO) which is an unavoidable overhead in communication systems. These offsets can be large when using moderate or low precision oscillators that drift due to temperature, humidity, and other intrinsic parameters. While today's IoT nodes use moderate-precision oscillators and independently adjust for CFO, a conductive substrate-enhanced IoT node can use a low-precision oscillator that is pulled to the correct frequency by an external high-precision oscillator that is connected via the conductive substrate, thereby obviating the need for each node to use more precise oscillators and separately perform CFO adjustment.

There are several other advantages of using conductive substrates for oscillator coupling. First, it is not complex to modify an IoT device to allow it to couple through a substrate since most MCUs and radios can use an external clock signal as reference, hence it is relatively easy to modify an IoT device to use an external low-precision oscillator. Second, oscillator coupling has low overhead compared to traditional clock synchronization methods since it does not incur protocol or compute overhead. Third, an IoT device only needs to leverage the substrate for oscillator pulling when needed. Many compute and sensing tasks only require a low precision clock, so the IoT device can judiciously use the the substrate to pull the oscillator only when communication is needed.

Contributions: In this paper, we present our method, **COCOON**¹ that makes the above idea practical i.e. we enable low-cost IoT nodes or RF elements to share a *virtual oscillator* via coupling over a conductive substrate (as shown in Figure 2).

¹COCOON = COnductive substrate-based COupled Oscillator Networks

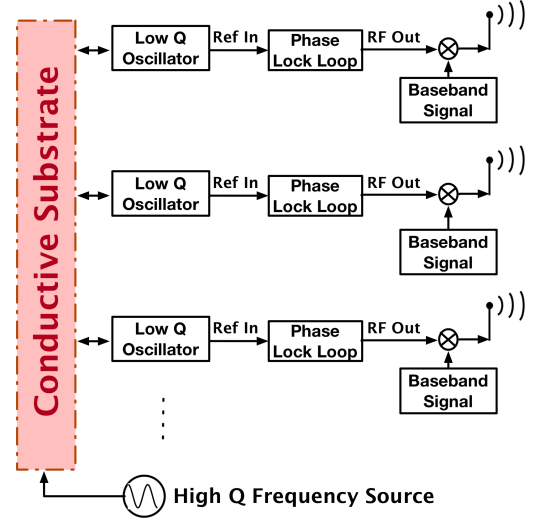


Figure 2: Overview of COCOON: Several low-Q oscillators as well as at least one high-Q oscillator are connected to a conductive substrate. The low Q oscillators feed into a Phase Lock Loop (PLL) which up-converts the input reference into an RF frequency for the radio.

Our work addresses several challenges to making this idea practical. First, while the general theory of oscillator coupling is well established, little is known about the effectiveness of coupling over different conductive substrates or the empirical performance of typical low-cost ceramic oscillators over such substrates. We show through empirical measurements that multiple ceramic oscillators on a conductive substrate can be pulled by a high-precision master oscillator at low power (less than 1mW).

Second, we address the question of how the high-precision oscillator can detect that it has successfully synchronized the low-precision oscillators. In the absence of such feedback, it not possible to design a closed-loop system that ensures synchronization at minimum power consumption. We design a novel passive measurement technique that measures the distance from quasi-lock state in a passive manner by just observing the frequency-domain subbands created by the oscillators. This requires no communication across IoT devices and the master oscillator making it easier to scale across devices that have different radios and use different protocols.

Third, we look at how pulling a low-frequency reference oscillator can be translated into RF frequency adjustments that can enable wireless communication. We show that we can indeed translate frequency pulling of a reference oscillator to RF frequency via a PLL. We show that **COCOON** can achieve distributed RF synchronization with low phase error of less than 0.08 radians at a fraction of cost of conventional solutions, thus providing a viable and low-cost method for pulling RF frequencies to minimize CFO.

Fourth, we demonstrate an end-to-end multi-static backscatter system that leverages **COCOON** for synchronization over a conductive substrate. We show that this system has high performance and can achieve nearly the same performance as a wired synchronization system that uses an Octo-Clock but that is 27× more expensive than **COCOON**.




Figure 3: Application of COCOON to multi-static backscatter to synchronize and cancel carrier interference.

2 APPLICATIONS OF COCOON

We now describe innovative applications enabled by COCOON in built settings like smart homes and warehouses.

Spatially Distributed Backscatter Readers: A key challenge in large-scale logistics is ensuring that all tags are accurately inventoried in a shipping warehouse (e.g. Alibaba). This is a challenging problem since tags are often blocked by other containers and there is significant multipath and attenuation. Even a small fraction of missed tags can have large revenue implications since logistics companies ship billions of packages a year.

Our work opens up an intriguing possibility that a network of RFID readers connected to a conductive substrate might be designed with low-cost ceramic oscillators as opposed to high-precision oscillators. COCOON can allow RFID readers to be spaced apart from each other while still allowing them to be tightly synchronized without requiring expensive cabling. Figure 3 shows such a multi-static backscatter setup using multiple synchronized RF carriers that illuminate the scene from different vantage points to achieve higher range and better spatial coverage.

COCOON has two key advantages over alternate methods: a) it is more than an order of magnitude cheaper than using dedicated wired synchronization modules (see Table 1), and b) it is more robust and more efficient than wireless synchronization methods [15, 22, 29] which can be unreliable in high multipath environments, use precious spectrum resources for wireless sync messages, and require complex DSP processing at each receiver.

Data center monitoring: Industrial IoT applications such as temperature and humidity monitoring in a data center involve acquisition of synchronous samples from distributed sensors in-order to estimate a spatio-temporal phenomena. One key challenge in these environments is that the frequency of clocks at the different nodes can vary significantly due to wide temperature variations across locations (frequency difference of clocks increases with temperature). This necessitates either frequent synchronization, which can be expensive in terms of resource usage, or the use of high-precision but roughly 50-100× higher priced temperature compensated oscillators (TCXO or OCXO). COCOON offers a much less expensive alternative since a single TCXO or OCXO can be the master oscillator that pulls a number of low-precision ceramic oscillators to compensate for temperature variations without requiring expensive oscillators or frequent wireless synchronization.

Distributed Beam-forming: An area of significant interest in 5G is supporting a large number of users and devices in challenging outdoor and indoor areas such as dense city squares, malls/offices, stadiums, train stations, factories and warehouses. COCOON enables a new way to distribute and synchronize nodes in a system,




Figure 4: Transmission line lump model for conductive substrate

thereby enabling better signals and data rates. Once nodes are synchronized through the wall-paper, COCOON can generate high frequency, long range and narrow beam lobes without requiring a centralized clock distribution or DSP processing unit for wire/wireless synchronization. Additionally, the scalability of COCOON architecture allows us to synchronize more nodes while keeping these individual nodes low power thus reducing the complexity of the individual transmitters.

Sharing Synchronization Information: Sharing an oscillator can be beneficial when a number of IoT devices that are coupled via a conductive substrate are communicating with a smartphone or access point that is not connected to the substrate. Rather than having to synchronize to each IoT device separately, a group of nodes that share an oscillator can couple via the conductive substrate. This may be particularly useful when devices are intermittently powered through energy harvesting [13] or highly duty-cycled and need to be synchronized upon wakeup.

3 PULLING IMPRECISE OSCILLATORS

COCOON builds on the idea of frequency coupling, a technique that has been long studied and applied in electronics design. The term injection locking is used to describe an oscillator that oscillates at the frequency which the oscillator is injected [31]. In addition to LC oscillators, this method has been shown to work for other oscillators including ring oscillators [24] crystal oscillators, and ceramic oscillators[28].

The two most important factors that affect coupling are the coupling factor and natural frequency difference between oscillators [14]. The famous Kuramoto model [26] models how these parameters affect a system of N inter-connected oscillators:

$$\frac{\partial \theta_i}{\partial t} = \omega_i + \frac{K}{N} \sum_{j=1}^N \sin \theta_j - \theta_i \quad (1)$$

In Kuramoto's model, coupling factor K is one of the factors that determines the state of locking. This constant factor represents how much an oscillator would receive the injected oscillation.

COCOON brings the idea of frequency pulling to low-precision oscillators connected to a high-precision oscillator over commonly available conductive substrates. For this to be feasible, we need to answer three questions: a) Are the intrinsic physical properties of conductive substrates conducive to frequency pulling, and if so, over what distances? b) Can frequency pulling over a conductive substrate work for frequency drifts observed in typical ceramic oscillators? and c) Do ceramic oscillators have the intrinsic ability to be pulled towards another frequency?

Suitability of conductive substrates: We now try to understand whether the properties of typical conductive substrates are




Figure 5: (a) Beat frequencies are observed once resistance increases above 1.3 kΩ (roughly tens of meters across conductive substrates), (b) Substantial differences of up to 40kHz is acceptable for frequency pulling, hence even very low precision/high drift oscillators can be pulled together.

conductive to oscillator pulling. To explore this, let us first look at how to simulate the behavior of a conductive surface based on the above model. For a conductive substrate, the series resistance and parasitic capacitance to the dielectric are two factors that affect how much energy is lost in the transmission and thus how much a receiver oscillator “feels” the injected signal [11].

Figure 4 shows a micro-electronic transmission line lump model that we use to represent our conductive substrate. Since a conductive substrate is not closely intertwined with a ground, low frequency transmission of 1MHz incurs negligible current leakage due to little capacitive coupling between the transmission line and ground (impedance Z from the substrate to ground of the dielectric capacitor $\frac{1}{j\omega C}$ is extremely high). Thus, resistance becomes the dominant factor in determining the coupling factor, K .

Based on the above model, we investigate frequency coupling using a SPICE simulation. We built an equivalent lump circuit model that can represent a realistic multi-oscillator connection with a conductive substrate. We vary the natural frequencies and resistances between oscillators to understand synchronization behavior.

Impact of coupling factor: One way to gauge the quality of synchronization between coupled oscillators is the presence of beat frequencies, thus our first investigation looks at the magnitude of beat frequency signal components for different amounts of coupling. Figure 5(a) shows the relation between the change in observed beat frequency and resistance between two mutually injected oscillators running at 1MHz. We see that there is no observable beat frequency under 1 kΩ and very small beat frequencies between 1 kΩ and 1.5 kΩ. As the resistance increases, two oscillators become de-coupled and the beat frequency approaches $|f_2 - f_1|$. A typical 1 m long conductive copper tape of 5cm width has 1 Ω/meter resistance; a single conductive thread has 50 Ω/meter resistance, and the conductive paint has 55 Ω/meter resistance. This means that two oscillators can synchronize over several tens of meters over typical conductive substrates which is more than sufficient for most practical home or warehouse deployments.

Impact of frequency difference: A key factor that determines the coupling state is the frequency difference between the oscillator and incoming oscillation. When two oscillators are coupled, it is known that the final settling frequency is a tone that is between the original frequencies [35]. Again, we use the presence of beat frequencies as a way to gauge the quality of oscillator synchronization; simulation results in Figure 5(b) shows how much frequency

difference can be tolerated before beats start to occur. We observe that if $|\delta f| < 40\text{kHz}$, then the oscillators converge to the same frequency. Since beat frequencies only starts to appear after about 40kHz, we can tolerate up to 20000 ppm drift between oscillators! A low-end quartz-based commercial 1MHz oscillator has a frequency deviation of 80ppm which means only 160Hz maximum frequency difference, which is only a fraction of the range. A lower cost but less precise ceramic resonator has an initial accuracy of 5000 ppm, and drifts significantly with temperature and age (2000 ppm) [34]. This is also within the regime for synchronization. A 9 stage ring oscillator working at 33.45MHz range has been reported to have an accuracy of around 13000 ppm [27]. Even this is potentially within the range that we can support. This opens up use cases where we mix-and-match precise but higher cost oscillators with imprecise-but-lower cost oscillators.

Impact of Q Factor: The pulling range i.e. the maximum difference between the slave and master frequencies as described above also depends on the “bendability” of the slave oscillator. This “bendability” is determined by the following equation [31]:

$$\omega_L \approx \frac{\omega_0 I_{inj}}{2Q I_{osc}} \quad (2)$$

where ω_0 is the free running frequency of the oscillator and Q is its Q factor, I_{inj} and I_{osc} are the injection current and oscillator output current, respectively. Note that with lower Q , the lock range is larger i.e. the oscillator’s frequency is more “bendable”.

COCOON leverages this fact by pairing low Q (more bendable) oscillators such as ceramic oscillators which have a Q factor of \sim several hundreds, while a high Q (less bendable) oscillator, i.e. quartz crystal has a Q factor of \sim tens of thousands. In our architecture, low Q , imprecise oscillator(s) will be pulled towards a high Q , precise oscillator which is what we would like to see. If a yet higher Q crystal oscillator is used, such as a Temperature-Controlled Crystal Oscillator (TCXO) or Oven-controlled Crystal Oscillator (OCXO) with Q factor of \sim hundreds of thousands, then the lower Q oscillator is pulled almost entirely towards the higher Q oscillator.

Scaling to large numbers of low Q oscillators: To demonstrate scalability, we modify our SPICE simulation to have a large number of low Q (set to Q factor=400) free running Hartley oscillators (up to 48), each outputting a frequency of 1MHz with 5000PPM at $\sim 100\mu\text{W}$. We then observe how much power it would take for a high Q oscillator oscillating at 985kHz to completely frequency pull all the low Q oscillators such that the resultant frequency of all oscillators is exactly 985kHz. Figure 6(a) shows that even with 48 oscillators, it would only take less than 150mW for the high Q oscillator to frequency pull the entire system of oscillators.

Scaling over large distances: We now ask whether the frequency pulling approach can work over long distances. According to Kuramoto’s model in eq.1, oscillators will injection lock with each other when the equivalent resistance is low enough. In our simulation, a pair of Hartley oscillator were frequency locked under a resistance of 1.5kΩ. Figure 6(b) shows the resistance per meter length at 10MHz of three different types of conductive substrates. We see that all these substrates have low resistance per meter length which implies that coupling should work over long distances. For example, the silver laminated wall paper has a resistance of 3 ohms




Figure 6: (a) Output power of high Q oscillator needed to frequency pull low Q oscillators, (b) Resistance of different types of conductive substrate at 1 meter.

per meter, which in theory means that we can synchronize over several hundred meters.

4 DETERMINING THE LOCKED STATE

Given that we can pull imprecise oscillators towards a precise oscillator to achieve synchronization, the next question is how can a master oscillator determine the synchronization state of the slave oscillators and adjust the frequency pulling power accordingly? This is particularly important if the master is battery-powered and wants to operate at the minimum power needed to achieve locking. **Sideband Lock Detection.** Let us start by looking at the frequency spectrum observed under injection locking. Injection locking is a nonlinear dynamic phenomenon[43]; when an oscillator is under injection, its output tracks the injected input frequency and the reactive components of the oscillator compensate for the instantaneous phase difference at the input. As a result, the oscillator frequency is pulled or locked towards the injecting frequency. In electronic oscillators, an oscillation-generating tank is used to provide a phase shift according to the injection phase difference. Simple circuit analysis on AC current with frequency ω through an ideal LC tank with feedback yields in the extreme, the transfer function of $\frac{V_{out}}{V_{in}}$ could have a $+\pi$ phase difference because the reactive components can lead or lag a sinusoidal wave. This maximum phase shift implies that there is an absolute limit of frequency offset that a particular oscillator can tolerate and lock on to; this frequency range is defined as ω_L . While within this locked range ω_L and also while under weak injection, beat frequencies will manifest in the frequency domain as side bands; this indicates the free-running oscillator is “pulled” more weakly by the injection frequencies. As the injection frequency becomes stronger, the oscillator is pulled more towards the injection signal and reduces the side bands, and with a large enough injection signal, side bands vanish indicating that the frequency is locked [31].

While the sidebands observed during injection locking provides a method for the master oscillator to track the level of synchronization of the slave oscillators, when should the master stop increasing the pulling power? Increasing pulling power until side bands disappear entirely is inefficient since it can take a significant amount of power to achieve this state.

Achieving Quasi-lock. A key observation we make is that it is not necessary to completely eliminate side bands to achieve lock. In other words, we do not need to increase power until all oscillators emit a pure tone. Rather, a frequency with majority power on a desired tone along with small power in side bands




Figure 7: Illustration of QL-Distance. If the distance is large, linearity is not preserved across frequencies and thus not considered quasi lock. The distance is positive if power is less than needed.

can also be acceptable. This is referred to as a *quasi-lock* state. In the time domain, quasi-lock can be observed as a slowly varying instantaneous phase; if the phase changes slowly enough, it will not impact communication.

Quasi-Lock Detection. To determine whether quasi-lock has occurred, we can use the fact that when in quasi-lock state, the magnitude of the side bands drops approximately linearly on a logarithmic scale due to an exponential decrease of the probability density in instantaneous frequency on a linear scale [31]. We define *QL-Distance* as a metric that captures the linearity of the side bands. To compute the QL-Distance, we identify the furthest observable side band peak ω_s and then form a straight line between the injection frequency and the furthest side band created. As shown in Figure 7, we compute the distance between the line and the furthest peak frequency with greatest magnitude, and refer to this distance as QL-Distance. A larger QL-Distance means that the side band spectrum is not very linear and unlikely to be in quasi-lock state. In our implementation, we compute the QL-Distance according to this method and empirically determine the threshold of QL-Distance to decide whether our system is in quasi-lock state.

Adaptive Power Control: We now look at how the master oscillator can leverage information about locking behavior to optimize system performance. Specifically, we focus on the case where the master is battery-powered and wants to determine the minimum power at which quasi-lock can be achieved.

The master starts by injecting a small amount of power to ensure weak injection. In this state the resulting frequency is not locked, as the side bands are large since the slave oscillators dominate. The master monitors the pulled oscillation in frequency domain, and determines the QL-Distance. When the output injection power keeps the subsequent side bands form a linearly decreasing pattern on the log scale (i.e. when the computed QL-Distance is close to zero), the master concludes that the system is in quasi lock state and keeps outputting such power as the minimum power needed to maintain synchronization for the system. If the master continues to increase its power, it will see that the injection frequency will have the majority of the power i.e. that QL-distance starts to increase in the other direction. Eventually, it will see an one tone injection frequency if the power is large enough.

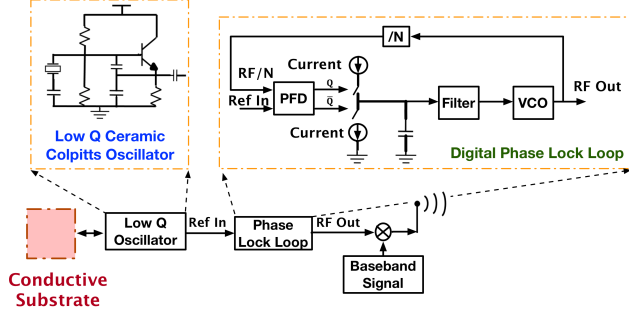


Figure 8: Block diagram of the building blocks of our system. A low Q ceramic Colpitts oscillator is used as the reference frequency that feeds into a Phase Lock Loop which generates the RF frequency for the radio.

5 PULLING THE RF CARRIER FREQUENCY

So far, we have focused on locking the low-frequency oscillator on each IoT device attached to the substrate but one of our goals illustrated in §2 is to be able adjust the RF carrier frequency on each device in-order to optimize communication. In other words, we wish to “remotely control” the RF carrier frequency by pulling its reference frequency via conductive substrate. This is an unconventional way of adjusting the RF carrier — normally, the reference frequency is fixed to that provided by a local crystal oscillator whereas in our case the reference itself is being pulled by an external source.

To make this possible, we need to ensure that changes in reference frequency are propagated through the Phase Lock Loop (PLL) that up-converts the reference frequency to RF frequency output in the radio. We propose to leverage the fact that the PLL is a feedback circuit that has a stability region within which it can function effectively. As long as the reference oscillator’s frequency is within this stability region, any changes in frequency of the reference should translate to a linear change in the RF frequency output. We start with a short background for how digital phase lock loop works and expand on our approach.

PLL Background: A Phase Lock Loop (PLL) is a feedback loop that provides stable, accurate, and clean oscillations. A common application for PLLs is generating RF carrier frequencies for radio communications. Figure 8 shows the basic functional blocks of a PLL. The RF signal is output from a voltage controlled oscillator (VCO) whose frequency is controlled by a voltage input provided by the output of a charge pump. In order for the RF carrier frequency to remain stable, the feedback loop consists of two parts: 1) A frequency divider using a counter to reduce the RF output to a lower frequency and 2) a comparator that compares the output of the divider to a reference source through XOR operation using a phase frequency detector; the charge pump’s output voltage is controlled by this phase frequency detector to adjust the carrier frequency, thus closing the loop.

In the case of COCOON, we steer the RF output frequency remotely by adjusting the reference frequency via the conductive substrate rather than using a VCO, as is the case with a conventional PLL. To get a sense of how a slight mismatch of reference frequency would steer the RF frequency, we take a closer look at the phase frequency detector.

Phase Frequency Detector: A phase frequency detector determines the inconsistency between two input waves, the reference frequency and the divided RF wave, and outputs a voltage as a control to the voltage control oscillator. In a digital phase lock loop, the inputs are digitized into square waves which are sent into the phase frequency detector that outputs digital 1s and 0s for the control of the charge pump (illustrated in Figure 8). The two outputs, Q and \bar{Q} , control the gate of positive or negative current. The duration of the charging current in Q and \bar{Q} charges up the the capacitor C_p to a different voltage level and such different voltage is used to control and adjust the VCO to output its final RF frequency.

Impact of perturbations in reference frequency: Let us now look at the effect of small changes in reference frequency on the RF frequency output. We are interested in the effect on the lock range of a PLL i.e. the reference frequency range that the PLL is able to follow and lock. This is the regime within which a slave oscillator’s reference frequency should be adjusted by a master oscillator.

The lock range is given by [1]: $\Delta\omega_l = N\pi\zeta\omega_n$ where ζ is the damping ratio of an nth order feedback system in PLL and ω_n is the natural running frequency of the VCO and N is a scaling integer depending on the type of PFD used. Once the reference frequency drifts out of this range, the PLL is not locked and thus no reliable RF frequency can be used.

6 ENABLING SYNCHRONIZED MULTI-STATIC RF BACKSCATTER

The ability to pull the output of a PLL using a remote reference presents several intriguing new opportunities for distributed communication systems that are inter-connected using flexible conductive substrates. In this work, we look at one such application, i.e. multi-static backscatter communication.

A multi-static backscatter system comprises multiple carrier transmitters that are spatially distributed and that concurrently illuminate a shared region to provide better spatial coverage and better signal-to-noise ratio. Multi-static configurations are typically either wired (e.g. Ethernet-connected readers such as in RFGO [6]) or wireless (e.g. PushID [40]); we offer an interesting alternative that involves distributed carrier transmitters connected over low-cost, flexible conductive material. Such a multi-static backscatter system is considerably cheaper and easier to install than traditional Ethernet-connected readers and can therefore enable more flexible configurations at equivalent performance but without the complexity of fully wireless configurations.

A key issue in multi-static backscatter systems is the need for tight synchronization. A traditional mono-static backscatter system such as a commercial RFID reader does not need to synchronize since the carrier transmitter and receiver are co-located at the same device and can share the same clock source. However, in a multi-static backscatter system, frequency synchronization is essential in-order to be able to effectively filter out the carrier signal and extract the weak backscattered signal from the tags. Frequency synchronization is particularly important when the clock sources can have large drift, which is the case when using low-cost clock sources such as ceramic oscillators.

To demonstrate how COCOON can enable high-performance multi-static backscattering, we designed a prototype comprising




Figure 9: Multi-static backscatter setup: Two carrier transmitters illuminate the scene and a single receiver cancels the carrier signal using a direct-down conversion mixer to extract the weak backscattered signal. The carrier transmitters are equipped with low-precision ceramic oscillators synchronized via a conductive substrate.

two transmitters and one receiver that integrate frequency pulling on silver laminated wall paper shown in Figure 9. The receiver uses a direct-down conversion mixer to convert the modulated RF signal into base-band and a DC blocker to eliminate direct carrier interference that is generated from two distributed transmitters.


7 IMPLEMENTATION

We implement **COCOON** on several conductive substrates using off-the-shelf components and analyze its wireless communication potential with a software-defined radio (USRP).

In a conventional radio front end, the reference source is typically a piezoelectric resonator (a quartz crystal). Such a resonator does not have any external feedback to allow for frequency injection. Instead, we use an active Colpitts RC oscillator that comprises an off-the-shelf low-Q ceramic resonator. A Colpitts oscillator has a feedback loop that can perform injection locking. We use standard RC components with the popular 2N2222 NPN transistor to build such oscillator. We implement the oscillator on a PCB board of size $2\text{ cm} \times 4\text{ cm}$ with a contact spring connector that easily connects to the conductive substrate (Figure 10(a)). The output frequency of our ceramic Colpitts oscillators are centered around 13.35 MHz. We chose the 13.35MHz oscillator because this frequency is an available reference input to common off-the-shelf PLLs. In addition, the resistance per meter across conductive substrate is low at this frequency as shown in Figure 6(a).

Implementation with conductive substrates: We use three different types of conductive substrates. The first is silver laminated paper stock that was sourced from a local paper company and is conventionally used in labels and signs [38]. We obtained a 300 ft reel and use a portion of it for evaluation. The second is copper conductive tape that is usually used for noise shielding for electric acoustic instruments (from a music hardware store). The third is conductive paint that turns ordinary surfaces into smart ones and is popular for interactive wall installations [37].

We implement **COCOON** on all three of these substrate on a wall, each with length of 5.5 meters and width between 2-3 cm. Different substrates under the same length have different impedance, which



(a) Ceramic Colpitts Oscillators

(b) Experiment setup

Figure 10: (a) Our PCB implementation of a Ceramic Colpitts Oscillator. (b) Experimental setup: Six oscillators are connected in a linear chain on a 5.5m long conductive laminated wall paper, conductive tape and conductive paint.

will exhibit different pulling behavior. We assemble our Colpitts oscillators onto the conductive substrate as shown in Figure 10(b). We place the high Q frequency generator at one end of the substrate, which can generate an output at 2dBm.

RF System Implementation: We ensure **COCOON** is at quasi-lock by outputting $>-2\text{dBm}$ from the frequency generator. We then vary the frequency output of the frequency generator and monitor both the resultant pulled frequency of our oscillators and the output RF frequency of our PLL using a spectrum analyzer.

The RF front ends of the transmitters and receiver comprise of three 13.33MHz ceramic oscillators attach to the conductive substrate which connects to controllable ADF4351 phase lock loops [2] to generate the RF carriers. The ADF4351 PLL consists of a local tunable VCO for RF generation, as well as a feedback loop with digital counter for frequency division. Our setup includes two transmitters transmitting 915MHz carrier tones. The carrier signals from the two transmitters are amplified and transmitted through antennas as multi-static carrier waves. The receiver uses a direct-down conversion mixer MAX2021 [23] to convert backscattered RF signal into base-band. The LO to the mixer is provided by a third PLL which attaches to the conductive substrate. We decode the base-band signal with low frequency USRP LFRX. A backscatter tag is placed 2.5m away from the transmitters and it is in the center of the room. The backscatter tag transmits ASK symbols.

We built 6 ceramic oscillators, and we exhaustively evaluated the communication performance for 20 combinations out of these at Tx and Rx. We also implemented the set up on 3 different dimensions of $1\text{ m} \times 0.1\text{ m}$, $3\text{ m} \times 0.1\text{ m}$ and $6\text{ m} \times 0.1\text{ m}$.

8 EVALUATION

We now turn to empirical experiments to validate our ideas in practice. While we have experimented with several substrates, we show a few key results on all substrates and in other cases only for the conductive wallpaper since trends are similar.

8.1 Coupling with conductive substrates

In this experiment, we look at frequency coupling through conductive wallpaper. We look at nodes connected via a linear topology which represents the worst-case for coupling behavior since the end-points are not directly connected but need to synchronize via “relays”. We assembled six 1MHz crystal Colpitts oscillators using off-the-shelf quartz crystals and electronic components, and spaced

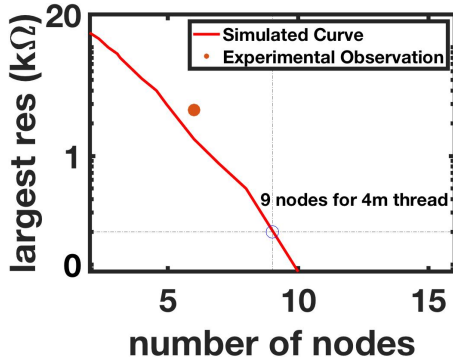


Figure 11: Effect of number of nodes in a linear topology on max. resistance between adjacent nodes for coupling (700 Hz freq. difference). Experimental results are close to simulations. Roughly 9 nodes can be connected using a 4m conductive thread, paint, or tape with resistance around 50 Ω /m.

these evenly on a 550cm \times 3cm conductive wallpaper. Figure 10(b) shows the experimental set up. We vary the resistance between the oscillators to understand whether our simulation model in §3 matches to empirical behavior.

Figure 11 compares our simulation versus experimental result in terms of coupling requirements between adjacent nodes as the number of nodes increases. In our experimental setup, we see that locking happens when the resistance is roughly 2.5k Ω (red dot in the figure). This is quite close to the number predicted by our simulations (line). Given that the actual resistance of the conductive sheet, paint and tape is roughly 50 Ω /m, we should be able to scale to even more nodes. From the simulation curve, we can estimate that roughly nine nodes can synchronize over a 4m length of conductive tape, paint or tape which has resistance around 50 Ω . This illustrates that deployments over moderately large conductive surfaces are feasible with COCOON.

8.2 Pulling low precision oscillators

Next, we validate our claim that low Q oscillators can be pulled easily by looking at the frequency pulling of low-cost ceramic oscillators. As discussed in §3, oscillators with lower Q value have a larger frequency pulling range. Such range allows the oscillators to be pulled towards an injection frequency. Due to low Q, a ceramic oscillator has a larger error margin that typically results in a highly inaccurate frequency. As a result, PLLs in radio front ends do not use ceramic oscillators as their reference input despite the fact that they are very inexpensive.

Averaging among ceramic oscillators: First, we look at coupling between a group of ceramic oscillators, all of which are transmitting their signal at the same power. If multiple ceramic oscillators at the same power are mutually coupled, the resultant output frequency is the average between them.

This can be leveraged in two ways. First, the group of ceramic oscillators are now synchronized with each other can perform tasks like synchronized sampling, for instance, to enable audio beamforming. Second, the nodes can wirelessly communicate with each other with no carrier frequency offset (as long as the resulting average frequency is within the lock range of the PLL as described

in §5). The averaging effect can also pull the frequency of outlier oscillators closer to the mean.

To illustrate, we mutually inject ceramic oscillators through our laminated wall paper. We measure the free running frequency of each oscillator, and then add it to the laminated wall paper one by one. We record the resulting frequency each time we add a new ceramic oscillator on board.




Figure 12: Ceramic oscillators are added one by one and the resulting averaging effect due to frequency pulling is seen. The bar is the free-running frequency of each oscillator and the line is the average frequency.

Figure 12 shows the results. We see that whenever a new oscillator is added to the conductive substrate, the resulting frequency is the average of the existing oscillator(s) and the new oscillator. For example, when a higher drift second oscillator is introduced, the first one pulls it towards itself; similarly, the fourth oscillator has high drift but is pulled towards others.

Pulling towards a high-Q, high power oscillator: If we now attach a high Q oscillator such as a fine crystal oscillator with sufficient injection power to the substrate, the ceramic oscillators bend entirely towards the frequency of the high-Q oscillator. We used a high-Q oscillator at 13.33MHz and attached it to the above substrate and all the ceramic oscillators locked to this frequency. Thus, we can finely control the resultant frequency by fine tuning a single high Q oscillator.

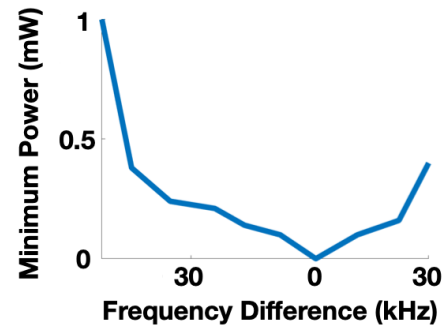


Figure 13: Min. power to quasi-lock for different frequency differences between master and slave oscillator.

8.3 Evaluation of Quasi Lock State Detection

We have shown that ceramic oscillators can be bent towards a desired frequency; we now turn to looking the properties of the




Figure 14: Min-power for Quasi Lock as number of ceramic oscillators increases. The power is less than 1mW even with four oscillators with frequency variation of 5000PPM.




Figure 15: QL-Distance is smallest when in quasi-lock state (15(a)) and is higher both when unlocked and when power exceeds quasi-lock (15(b)).

quasi-lock state and validate our claim that the quasi-lock state can be detected by looking at the linearity of the sidebands. We also look at how much power benefits are obtained by using quasi-lock as the stopping criteria versus waiting for full locking.

Illustration of Quasi Lock. First, we illustrate the side band structure in three stages — a bit before quasi-lock, in quasi-lock and a bit beyond quasi-lock. This is what a master oscillator will observe as it slowly ramps up the power that it feeds into the substrate.

We show this with a ceramic oscillator being pulled to a waveform generator's signal as the waveform generator slowly increases its power. The waveform generator can be considered an unbendable oscillator, hence this represents the case of a low Q and high Q oscillators being connected together. Using spectrum analyzer, we capture these changes on the substrate and plot the results.

Initially, we do not get frequency locking if the output power is too low (QL-Distance is around 10). Shown in figure 15, as input power increases, we see quasi locking when the output power of the high precision source is large enough (QL-Distance is around 5). As power is increased further, the master oscillator frequency becomes more dominant and the linearity of side bands diminishes (magnitude of QL-Distance starts increasing back to 10). Eventually, the master oscillator frequency dominates and the side bands diminish.

In our experiments, we find that the quasi-lock condition occurs when QL-distance ≤ 6 (across all substrates). Hence, we use QL-distance = 5 as our threshold in our experiments. We increase our high Q oscillator power until the QL-distance is smaller or equal to the threshold. The power required to pull thus becomes the minimum power from the high Q oscillator to ensure quasi-lock.




Figure 16: Power required for quasi lock vs. full lock for three oscillators on different substrates. There is substantial gains from using quasi-lock as the stopping condition.

Minimum power for quasi lock: Next, we look at how much power needs to be injected to achieve quasi-lock state. Figure 13 shows the effect of frequency difference between a low-Q and high-Q oscillator on power consumption. The larger the frequency difference, the larger the power it takes for the system to attain quasi lock. However, the power required is quite small for typical range of frequency variation. For example, the frequency variation one can expect in a 5000ppm low-precision ceramic oscillator is less than 65 kHz and we can see from the figure that it takes about 1 mW of power to achieve quasi-lock in this case. Hence, we can see that quasi-locking can be achieved at low power.

Figure 14 shows the power consumed as the number of ceramic oscillators increases from one to four. We see that power increases as the number of oscillators grows but still we only need 1 mW to lock four oscillators.

Power needed for quasi-lock versus full lock: We now turn to a comparison of the power needed for synchronizing a set of ceramic oscillators to the quasi-lock state versus fully locked state.

Figure 16 shows results for synchronizing three ceramic oscillators with a 20kHz frequency difference from the high Q source on different conductive substrates. We see that there is significant difference between the power consumed for quasi-locking versus locking. There is about a $4.5\times$ difference between the power needed for achieving these two states for the conductive wallpaper, a $7\times$ difference for the conductive tape, and more than $11\times$ difference for the conductive paint. This shows that there is significant power benefits to using the quasi-lock criteria for pulling oscillators.

Note that the above results are for the power consumed by the master oscillator. The power consumed at the slave ceramic oscillator at 4V is only $\sim 90\mu\text{W}$ which is very low.

8.4 Performance of Adaptive Power Control

We now look at the adaptive power control method that is used by the master oscillator to determine whether the oscillators attached to the substrate are in quasi-lock condition.

Adaptive power control: To demonstrate the method, we start with a single ceramic oscillator on the substrate, then add a second oscillator, and then a third. Figure 17 shows how the QL-Distance observed by the master node varies as oscillators are added. Initially, the system is in quasi-lock state and the QL-distance is less than the quasi-lock threshold (5 in our case). When the second oscillator is added, the QL-Distance increases to about 12. The master starts

ramping up the output power until it sees that the QL-Distance drops below the quasi-lock threshold again. This repeats two more times as additional oscillators are added to the system.




Figure 17: Adaptive power control as oscillators are added one-by-one. The master oscillator ramps up power when it observes QL-Distance to be higher than threshold until it reaches quasi-lock.

QL-Distance trend: We now look at the trend in QL-Distance in a bit more detail. We ask whether the QL-Distance decreases monotonically as power increases; if not, the master oscillator cannot easily track whether the trend is reducing.

Figure 18 shows how the QL-Distance varies as input power is increased for conductive wallpaper. We see that the trend is generally quite linear and the QL-Distance decreases quite predictably as power increases.

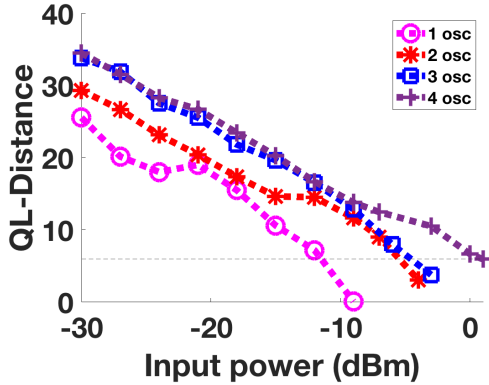


Figure 18: As injection power of the master oscillator increases, QL-distance reduces almost linearly. More oscillators requires more power to quasi lock, hence the curves shift towards the right.

8.5 Pulling RF output

So far, we have looked at pulling the frequency of the reference clock to the PLL. In this section we demonstrate that pulling the reference frequency in COCOON can in-turn reduce Carrier Frequency Offset (CFO) in the output RF frequency of the PLL.

Synchronization accuracy The accuracy of synchronization can be determined by looking at the misalignment of the phase between any two nodes. Phase misalignment can be estimated by looking at the difference of instantaneous frequency for a short period of time as shown in the following equation [45]:

$$\phi(t) = (F_1 - F_2) * t_{measure} + \phi(t_0) \quad (3)$$

Where $F_1 - F_2$ is the instantaneous frequency difference between two nodes, $t_{measure}$ is the measurement time period and $\phi(t_0)$ is the initial phase offset. After synchronizing two PLLs through the conductive wall paper, we measure the instantaneous frequency difference $F_1 - F_2$ across 4μ seconds (which is the symbol duration in 802.11[29]).




Figure 19: 95 percentile is under 0.1 radians for COCOON and more than 0.6 radians without it.

Figure 19 shows the distribution of the phase error. At low power communication whereas the SNR is under 10db, COCOON can ensure the degradation of SNR due to phase error is no more than 0.1db [29]. This is well within the margin of acceptable degradation — in narrow-band channels with Gaussian noise, a 0.1db SNR degradation results in less than 10% change in BER [16, 17].

Controlling RF output through reference input In §4, we had described that there is small margin of reference input frequency that remains within the lock range of a PLL, and such margin affects the output frequency approximately linearly. We empirically look at this range in the case of our PLL implementation.

Figure 20 shows the reference frequency (left y-axis) and the corresponding RF frequency of the PLL output (right y-axis). Note that the reference frequency corresponds to the output of a ceramic oscillator that is being pulled by a master oscillator. We see that the PLL changes linearly with the reference oscillator within lock range but when the reference oscillator is outside these bounds, the PLL does not lock. Within this range, we see that changing the reference frequency by about ~150kHz linearly changes the RF output of the PLL by about 2MHz. Thus, we see that COCOON can finely control the output RF frequency even if the nodes are using an inaccurate ceramic oscillator as reference input.

8.6 Performance of Multi-static RF Backscatter

We now examine how COCOON can offset the inherent high drift in cheap ceramic resonators to create a high-performance multi-static backscatter system. While the advantage of multi-static backscatter over mono-static backscatter has been established in prior work [20, 25, 39], the design of a multi-static backscatter system is considerably more complex. In this work, we focus on how our low-cost synchronization can improve the performance of such a multi-static backscatter system. We demonstrate this in three

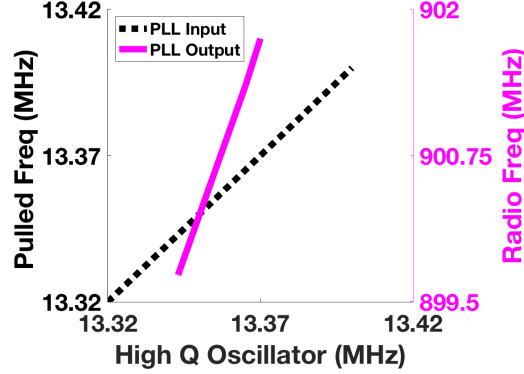


Figure 20: Dotted line shows the pulled reference frequency and solid line shows the RF frequency output of PLL. The PLL only works in part of the range of the reference; within this range, there is a linear relation.

ways — first, we show that the received BER with COCOON is considerably more stable than if we used unsynchronized oscillators; second, we show that frequency pulling achieves equivalent performance of high-quality sync using an Octo-Clock but at a fraction of the cost; and third, we show that frequency pulling is more scalable than using a single source to feed a reference frequency to all nodes on the substrate.

Frequency pulling vs. Unsynchronized operation: Figure 21 compares the performance of frequency pulling versus unsynchronized operation. 20 combinations of 6 different ceramic oscillators provided to the Tx and Rx with and without using frequency pulling through the wall paper. We saw that for all combinations of ceramic oscillators, COCOON has nearly zero BER whereas the case where ceramic oscillators supply reference to the PLL independently has significant BER and poor performance.

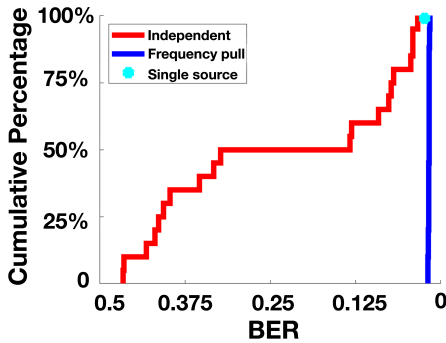


Figure 21: CDF of BER with and without frequency pulling. Frequency pulling yields a much smaller range for the same set of ceramic oscillators

Frequency pulling vs Octo-Clock: We also compare COCOON against a high-end synchronization solution i.e. use of an Octo-Clock from National Instruments. The Octo-Clock is commonly used in multi-static backscatter configurations [6] since it provides high-quality synchronization across nodes.

Figure 22 compares the performance when we synchronize the PLLs with the Octo-Clock without the wall-paper versus using frequency pulling over the wallpaper. We see that at -2dBm pulling

	COCOON	Octo-Clock
Oscillator PCB	2 x \$5	n/a
Wall Paper	\$10	n/a
Transmission lines	n/a	\$50
Master Waveform Generator	\$30	\$1333
Total	\$50	\$1383

Table 1: COCOON offers significant advantages in terms of price and ease of deployment as compared to a traditional Octo-Clock approach for synchronization.

	Vdvt [45]	MMIMO [29]	COCOON
Wire or wireless	Wire	Wireless	Wire
Core hardware used	NI-5791	USRP	Oscillator
Hardware cost	\$13k	>\$1.5k	\$50
DSP Required?	No	Yes	No
phase error(95%)	0.02	0.05	0.08
Signal degradation	<0.1db	<0.1db	<0.1db

Table 2: COCOON is an 1-2 orders of magnitude less expensive than alternate methods while having marginally higher phase error (95 percentile phase error is <0.06 radians larger from Vidut and 0.03 more than MegaMIMO).

power, the BER with frequency pulling is only $\frac{1}{5}$ higher than Octo-Clock based synchronization whereas the BER with a single source is $\frac{1}{2}$ higher than Octo-Clock.




Figure 22: At -2dBm pulling power, the single source BER is 0.002 high than frequency pulling while the high end Octo-Clock is 0.001 lower.

This result is particularly important when one considers the monetary cost of these approaches. Table 1 compares the cost of purchasing an Octo-Clock versus COCOON with ceramic oscillators. We see that the cost of COCOON is roughly 27× lower while offering comparable performance.

Frequency pulling vs. Single source: We look at how COCOON performs against a system that uses a single frequency source, which distributes using the conductive substrate.

Figure 23 shows the signal to noise ratio of the received signal between single source and frequency pulling at three different dimensions. We see that as the wall paper increases in length, frequency pulling performs better than having a single source. This is because a clock signal attenuates when transmitted over a distance and varies due to noise and distortion on the conductive substrate both of which hurts the PLL output. COCOON avoids this by using an oscillator connected directly to the PLL's reference and pulling the oscillator rather than feeding a remote signal to the PLL.




Figure 23: SNR for single source is around 5dB at 1m and 3m, but it drops to 1 db at 6m. Frequency pulling is able to maintain above 5dB SNR for all three lengths.

Qualitative Comparison against Related Efforts: We now provide a qualitative comparison of the received signal performance, implementation cost, and hardware/software overhead across other wired and wireless synchronization methods. The comparison is shown in Table 2.

Wireless synchronization methods such as MegaMIMO[29] requires complex digital signal processing such as channel estimation and link layer computation which is usually performed on dedicated ASICs or power-hungry FPGAs or GPUs. This makes it a high-cost and relatively complex solution. Wired solutions do not incur the computational overhead but often require expensive specialized hardware components. For example, Vidyut[45] synchronizes reference frequency through power-lines and offers exceptional accuracy but it uses dedicated high-precision PLLs to accomplish this. **COCOON** requires no additional hardware and directly synchronizes oscillators on the nodes.

9 RELATED WORK

While there has been extensive work on oscillator coupling, conductive substrates and clock synchronization, our work is unique in that we look at the intersection between these subjects.

Oscillator coupling in electronic circuits: While oscillator coupling has been an extensively explored in the context of electronic oscillators, it has not been studied in the context of conductive substrates. Prior experimental work in this field has largely focused on the use of coupling at extremely small scales such as within a single IC chip [9, 10, 19, 32] and for microwave frequency and above (>3GHz)[3, 33, 44]. Our work focuses on the effectiveness of such coupling over larger scales (conductive substrates on walls), lower frequencies (13MHz) and IoT devices. **COCOON** also does not require integrated hardware design and can be an add-on to existence communication or computing nodes with minimal hardware modification.

Conductive substrates: Conductive substrates have largely been explored for sensing [8, 46] or as a large capacitive touch interface [41, 42]. Conductive paint, for example, is generally advertised only as being intended for applications with DC circuits at low voltages [12]. This paper take an in-depth look at new possibilities these substrates offer for connecting and optimizing IoT devices.

Clock Synchronization: There is a substantial body of work on CFO estimation and time synchronization over wireless links

[15, 22, 30]. These are very different from our efforts to use a conductive substrate for oscillator pulling. Unlike wireless synchronization, **COCOON** does not require a link or MAC layer protocol and therefore does not incur additional hardware complexity or protocol-layer overheads. There is also prior work on wired synchronization through power lines [29] which synchronize by tracking changes in the incoming frequency. **COCOON** uses injection locking rather than explicit tracking and synchronization, hence, it is both simpler and allows more fine-grained control over the resultant frequency. A few works targeted node synchronization in IoT environment [21, 36], which is closely related to **COCOON**'s application in terms of addressing the challenge of tidely synchronize low cost, low quality IoT nodes. These works present solutions in the application layer which usually involve in setting up servers and clients as well as establishing novel protocols. **COCOON** addresses the challenge in the hardware level which not only can achieve time synchronization without additional protocol overheads, but also enables frequency synchronization.

10 CONCLUSIONS

In this paper, we argue that conductive substrates can present a promising new opportunity for re-thinking the design of distributed IoT systems in built environments like homes and warehouses. We showed that oscillator coupling via conductive substrates can allow us to design IoT devices using low precision oscillators and yet pull them to behave like high-precision oscillators. We showed that such a system can work at low power and low complexity and requires no co-ordination between the different devices. We showed that such “remote” frequency pulling of a low-precision reference oscillator can also be used to pull the PLL output in a radio, thereby enabling a high-performance multi-static backscatter system out of low-precision components. Our solution opens up new directions in thinking about the use of conductive substrates, and can potentially open up many applications that need radio transceiver synchronization or sensor sampling synchronization in a low-overhead manner. Any network that is deployed over large conductive surfaces such as wallpaper, curtains, and tents can potentially leverage this technique to improve system performance, save communication bandwidth, or reduce manufacturing cost.

11 ACKNOWLEDGMENTS

We sincerely thank reviewers for their insightful feedback. The research reported in this paper was sponsored in part by the CCDC Army Research Laboratory (ARL) under Cooperative Agreement W911NF-17-2-0196 (ARL IoBT CRA), by National Science Foundation under Grant No. 1763524 and supported by the HAZEN paper company. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the ARL, NSF or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

REFERENCES

- [1] Phillip w Allen. 2021. DIGITAL PHASE LOCKED LOOPS (DPLL). [http://pallen.ece.gatech.edu/Academic/ECE_6440/Summer_2003/L070-DPLL\(2UP\).pdf](http://pallen.ece.gatech.edu/Academic/ECE_6440/Summer_2003/L070-DPLL(2UP).pdf)
- [2] Analog Devices [n.d.]. *Wideband Synthesizer with Integrated VCO*. Analog Devices. Rev. A.
- [3] G. Anzalone, E. Monaco, G. Albasini, S. Erba, and A. Mazzanti. 2016. A 0.2–11.7GHz, high accuracy injection-locking multi-phase generation with mixed analog/digital calibration loops in 28nm FDSOI CMOS. In *ESSCIRC Conference 2016: 42nd European Solid-State Circuits Conference*. 335–338. <https://doi.org/10.1109/ESSCIRC.2016.7598310>
- [4] S Syed Azim, A Satheesh, KK Ramu, S Ramu, and G Venkatachari. 2006. Studies on graphite based conductive paint coatings. *Progress in Organic Coatings* 55, 1 (2006), 1–4.
- [5] D. A. Badillo and S. Kiaei. 2004. Comparison of contemporary CMOS ring oscillators. In *2004 IEEE Radio Frequency Integrated Circuits (RFIC) Systems. Digest of Papers*. 281–284. <https://doi.org/10.1109/RFIC.2004.1320596>
- [6] Carlos Bocanegra, Mohammad A. (Amir) Khojastepour, Mustafa Y. Arslan, Eugene Chai, Sampath Rangarajan, and Kaushik R. Chowdhury. 2020. RFGo: A Seamless Self-Checkout System for Apparel Stores Using RFID. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking (London, United Kingdom) (MobiCom '20)*. Association for Computing Machinery, New York, NY, USA, Article 54, 14 pages. <https://doi.org/10.1145/3372224.3419211>
- [7] Leah Buechley, Mike Eisenberg, Jaime Catchen, and Ali Crockett. 2008. The LilyPad Arduino: using computational textiles to investigate engagement, aesthetics, and diversity in computer science education. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 423–432.
- [8] Leah Buechley, David Mellis, Hannah Perner-Wilson, Emily Lovell, and Bonifaz Kaufmann. 2010. Living wall: programmable wallpaper for interactive spaces. In *Proceedings of the 18th ACM international conference on Multimedia*. 1401–1402.
- [9] Wei L Chan and John R Long. 2008. A 56–65 GHz injection-locked frequency tripler with quadrature outputs in 90-nm CMOS. *IEEE Journal of Solid-State Circuits* 43, 12 (2008), 2739–2746.
- [10] Yuen-Hui Chee, Ali M Niknejad, and Jan M Rabaey. 2006. An ultra-low-power injection locked transmitter for wireless sensor networks. *IEEE Journal of Solid-State Circuits* 41, 8 (2006), 1740–1748.
- [11] Albert Chen. 2013. “Why Signal Always Be Loss in a High Speed High Frequency Transmission Line”. IPC APEX EXPO Conference Proceedings.
- [12] Bare Conductive. [n.d.]. What is Electric Paint: the composition and application of conductive paints. <https://www.bareconductive.com/news/what-is-electric-paint-the-composition-and-application-of-conductive-paints/>
- [13] Jasper de Winkel, Carlo Delle Donne, Kasim Sinan Yildirim, Przemysław Pawelczak, and Josiah Hester. 2020. Reliable Timekeeping for Intermittent Computing. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. 53–67.
- [14] Florian Dörfler and Francesco Bullo. 2014. Synchronization in complex networks of phase oscillators: A survey. *Automatica* 50, 6 (2014), 1539–1564.
- [15] Jeremy Elson and Deborah Estrin. 2001. Time synchronization for wireless sensor networks. In *Proceedings 15th International Parallel and Distributed Processing Symposium. IPDPS 2001*. IEEE, 1965–1970.
- [16] Andrea Goldsmith. 2005. *Wireless Communications*. Cambridge University Press, USA.
- [17] Daniel Halperin, Wenjun Hu, Anmol Sheth, and David Wetherall. 2010. Predictable 802.11 Packet Delivery from Wireless Channel Measurements. *SIGCOMM Comput. Commun. Rev.* 40, 4 (Aug. 2010), 159–170. <https://doi.org/10.1145/1851275.1851203>
- [18] Takashi Ichinomiya. 2004. Frequency synchronization in a random oscillator network. *Physical Review E* 70, 2 (2004), 026116.
- [19] Shehzaad Kaka, Matthew R Pufall, William H Rippard, Thomas J Silva, Stephen E Russek, and Jordan A Katine. 2005. Mutual phase-locking of microwave spin torque nano-oscillators. *Nature* 437, 7057 (2005), 389–392.
- [20] John Kimionis, Aggelos Bletsas, and John N. Sahalos. 2013. Bistatic backscatter radio for power-limited sensor networks. In *2013 IEEE Global Communications Conference (GLOBECOM)*. 353–358. <https://doi.org/10.1109/GLOCOM.2013.6831096>
- [21] Sathya Kumaran Mani, Ramakrishnan Durairajan, Paul Barford, and Joel Sommers. 2018. A System for Clock Synchronization in an Internet of Things. arXiv:1806.02474 [cs.NI]
- [22] Miklós Maróti, Branislav Kusy, Gyula Simon, and Ákos Lédeczi. 2004. The flooding time synchronization protocol. In *Proceedings of the 2nd international conference on Embedded networked sensor systems*. 39–49.
- [23] Maxim Integrated [n.d.]. *Mixer*. Maxim Integrated. Rev. A.
- [24] Behzad Mesgarzadeh and Atila Alvandpour. 2005. A study of injection locking in ring oscillators. In *2005 IEEE International Symposium on Circuits and Systems*. IEEE, 5465–5468.
- [25] Zeinab Mhanna and Alain Sibille. 2014. Monostatic Vs. Bistatic Backscattering gain for UWB RFID Real Time Location Systems Z.Mhanna. In *COST IC1004*. Krakow, Poland. <https://hal-imt.archives-ouvertes.fr/hal-01188065>
- [26] Yamin Moreno and Amalio F Pacheco. 2004. Synchronization of Kuramoto oscillators in scale-free networks. *EPL (Europhysics Letters)* 68, 4 (2004), 603.
- [27] Yang Pan. 2008. Ring Oscillator Frequency Measurements Using an Automated Parametric Test System.
- [28] A. K. Poddar and U. L. Rohde. 2012. Techniques minimize the phase noise in crystal oscillator circuits. In *2012 IEEE International Frequency Control Symposium Proceedings*. 1–7.
- [29] Hariharan Rahul, Swarn Kumar, and Dina Katabi. 2014. JMB: Scaling Wireless Capacity with User Demands. *Commun. ACM* 57, 7 (July 2014), 97–106. <https://doi.org/10.1145/2618413>
- [30] Prakash Ranganathan and Kendall Nygard. 2010. Time synchronization in wireless sensor networks: a survey. *International journal of ubicomp* 1, 2 (2010), 92–102.
- [31] Behzad Razavi. 2003. A study of injection pulling and locking in oscillators. In *Proceedings of the IEEE 2003 Custom Integrated Circuits Conference*, 2003. IEEE, 305–312.
- [32] Behzad Razavi. 2012. *Design of integrated circuits for optical communications*. John Wiley & Sons.
- [33] D. Shin and K. Koh. 2019. 24-GHz Injection-Locked Frequency Tripler With Third-Harmonic Quadrature Phase Generator. *IEEE Transactions on Circuits and Systems I: Regular Papers* 66, 8 (2019), 2898–2906. <https://doi.org/10.1109/TCSI.2019.2912422>
- [34] Silabs [n.d.]. Silabs The Pros and Cons of Consolidating Frequency Sources Using Oscillators and Clock Generators. <https://www.silabs.com/documents/public/white-papers/>
- [35] Steven Strogatz. 2018. *Nonlinear Dynamics and Chaos* (2 ed.). CRC Press, 6000 Broken Sound Parkway NW, Suite 300.
- [36] Bharath Sundararaman, Ugo Buy, and Ajay Kshemkalyani. 2005. Clock synchronization for wireless sensor networks: A survey. *Ad Hoc Networks* 3 (05 2005), 281–323. <https://doi.org/10.1016/j.adhoc.2005.01.002>
- [37] <https://www.bareconductive.com/>. [n.d.]. Bare Conductive - Electrically Conductive Paint.
- [38] <https://www.hazen.com/>. [n.d.]. Hazen - Laminated Wall Paper.
- [39] Nguyen Van Huynh, Dinh Thai Hoang, Xiao Lu, Dusit Niyato, Ping Wang, and Dong In Kim. 2018. Ambient Backscatter Communications: A Contemporary Survey. *IEEE Communications Surveys Tutorials* 20, 4 (2018), 2889–2922. <https://doi.org/10.1109/COMST.2018.2841964>
- [40] Jingxian Wang, Junbo Zhang, Rajarshi Saha, Haojian Jin, and Swarn Kumar. 2019. Pushing the range limits of commercial passive rfids. In *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI})*. 301–316.
- [41] Yuntao Wang, Jianyu Zhou, Hanchuan Li, Tengxiang Zhang, Minxuan Gao, Zhuolin Cheng, Chun Yu, Shwetak Patel, and Yuanchun Shi. 2019. Flextouch: Enabling large-scale interaction sensing beyond touchscreens using flexible and conductive materials. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 3 (2019), 1–20.
- [42] Michael Wessely, Ticha Sethapakdi, Carlos Castillo, Jackson C Snowden, Ollie Hanton, Isabel PS Qamar, Mike Fraser, Anne Roudaut, and Stefanie Mueller. 2020. Sprayable User Interfaces: Prototyping Large-Scale Interactive Surfaces with Sensors and Displays. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [43] Xiaolue Lai and J. Roychowdhury. 2004. Capturing oscillator injection locking via nonlinear phase-domain macromodels. *IEEE Transactions on Microwave Theory and Techniques* 52, 9 (2004), 2251–2261. <https://doi.org/10.1109/TMTT.2004.834579>
- [44] J. Xu, J. Hu, B. Ciftcioglu, and H. Wu. 2013. A 4–15-GHz ring oscillator based injection-locked frequency multiplier with built-in harmonic generation. In *Proceedings of the IEEE 2013 Custom Integrated Circuits Conference*. 1–4. <https://doi.org/10.1109/CICC.2013.6658554>
- [45] Vivek Yenamandra and Kannan Srinivasan. 2014. Vidyut: Exploiting Power Line Infrastructure for Enterprise Wireless Networks. In *Proceedings of the 2014 ACM Conference on SIGCOMM (Chicago, Illinois, USA) (SIGCOMM '14)*. Association for Computing Machinery, New York, NY, USA, 595–606. <https://doi.org/10.1145/2619239.2626329>
- [46] Yang Zhang, Chouchang Yang, Scott E Hudson, Chris Harrison, and Alanson Sample. 2018. Wall++ Room-Scale Interactive and Context-Aware Sensing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–15.