

Liquid: A Wireless Liquid Identifier

Ashutosh Dhekne
University of Illinois at
Urbana-Champaign

Mahanth Gowda
Penn State University

Yixuan Zhao
University of Illinois at
Urbana-Champaign

Haitham Hassanieh
University of Illinois at
Urbana-Champaign

Romit Roy Choudhury
University of Illinois at
Urbana-Champaign

ABSTRACT

This paper shows the feasibility of identifying liquids by shining ultra-wideband (UWB) wireless signals through them. The core opportunity arises from the fact that wireless signals experience distinct slow-down and attenuation when passing through a liquid, manifesting in the phase, strength, and propagation delay of the outgoing signal. While this intuition is simple, building a robust system entails numerous challenges, including (1) pico-second scale *time of flight* estimation, (2) coping with *integer ambiguity* due to phase wraps, (3) pollution from hardware noise and multipath, and (4) compensating for the liquid-container's impact on the measurements. We address these challenges through multiple stages of signal processing without relying on any feature extraction or machine learning. Instead, we model the behavior of radio signals inside liquids (using principles of physics), and estimate the liquid's *permittivity*, which in turn identifies the liquid. Experiments across 33 different liquids (spread over the whole permittivity spectrum) show median permittivity error of 9%. This implies that coke can be discriminated from diet coke or pepsi, whole milk from 2% milk, and distilled water from saline water. Our end system, *Liquid*, is cheap, non-invasive, and amenable to real-world applications.

CCS CONCEPTS

• **Computer systems organization** → **Sensors and actuators**; • **Applied computing** → *Physics*;

ACM Reference Format:

Ashutosh Dhekne, Mahanth Gowda, Yixuan Zhao, Haitham Hassanieh, and Romit Roy Choudhury. 2018. Liquid: A Wireless Liquid Identifier. In *MobiSys '18: The 16th Annual International Conference on Mobile Systems, Applications, and Services, June 10–15, 2018, Munich, Germany*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3210240.3210345>

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.

MobiSys '18, June 10–15, 2018, Munich, Germany
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5720-3/18/06...\$15.00
<https://doi.org/10.1145/3210240.3210345>

1 INTRODUCTION

Liquids are known to be identifiable by their *permittivity*. At a high level, permittivity is the resistance that a liquid offers when an electrical field is formed inside it. Permittivity, however, is very difficult to measure. Today, devices that estimate the permittivity of a liquid rely on optical spectroscopy. These are large bulky equipment that cost in the excess of \$40,000 and hence are limited to laboratory settings [23]. Smaller devices exist at \approx \$18,000, but sacrifice accuracy and require dipping a probe into the liquid [33]. In this paper, we ask *whether it is feasible to estimate permittivity and identify liquids by non-invasively shining wireless signals from cheap off-the-self radios*. If successful, the applications could be many, ranging from airport security on liquids, quality control on stored drinks/wines in warehouses, cheap detection of water contamination especially in countries with limited access to clean water, all the way to futuristic cups that can measure the calorie count of a drink inside the cup, and devices that can analyze the blood content. This paper takes the initial small step towards this vision by developing the core capability first i.e., a reliable wireless liquid identifier.

We present *Liquid*, a cheap lightweight system that uses wireless signals to identify liquids. Liquid uses two cheap off-the-self ultra-wideband (UWB) radios [18]. It shines an UWB signal from one side of a liquid container and receives the signal on the other side. Liquid then analyzes the characteristics of the received signals to estimate the permittivity and identify the liquid. In designing Liquid, we build on several past papers that use wireless signals to classify liquids [25, 28, 35, 39]. The most recent of these papers [39] demonstrates the ability to classify 10 liquids using RFID readers and tags. However, the solution applies only to a specific set of 10 liquids that exhibit special wavelength relationships to each other (see Section 6 for more details). In contrast, Liquid estimates the intrinsic permittivity of any given liquid. Our experiments span over 33 different liquids, essentially demonstrating a general solution to liquid identification.

Liquid uses UWB signals with a 4GHz center frequency and a 1GHz bandwidth to measure two key attributes of liquids: *refractive index* and *loss factor*. These measurable attributes ultimately map to permittivity. When the UWB signal passes through the liquid, the liquid interacts with the signal in the two following ways:

(1) It slows down the signal's propagation speed. The slow down is referred to as *refractive index* which in turn reduces the wavelength of the signal (since speed of propagation $v = f\lambda$). This manifests as a change in the phase of the received signal.

(2) The liquid also attenuates the signal's amplitude. The attenuation is called *loss factor* and translates to a weaker received signal strength (RSSI).

By measuring the phase and the RSSI of the signal at the receiver, we had expected to compute the above two attributes. Unfortunately, since the phase wraps multiple times (around 2π) inside the liquid, the measured phase is not a unique signature¹. However, if the absolute propagation delay can also be estimated alongside the phase, perhaps the wrapped-phase can be unwrapped. Now, given that wireless propagation delay is in the granularity of picoseconds, precise estimation is a challenge with today's UWB receivers. Even clock synchronization is difficult to achieve to this consistent precision [18]. Moreover, noise, multipath, and the plastic material (of the container), all affect the phase and RSSI, polluting the measurements drastically. Extracting out the signal slowdown and the attenuation, in face of all these complications, forms the root of challenges in liquid identification.

LiquidID addresses these challenges through successive stages of signal analysis and engineering. Two key ideas underpin the solution.

(1) By using a simple wire connected between the transmitter and receiver, *LiquidID* develops a reference frame for time. In other words, the wireless transmission is forked between the wire and the wireless antenna, so that the receiver can use the wired signal as the baseline. This obviates the need for clock-synchronization, since measurements through the liquid, as well as through the air, can both be normalized over the same denominator. Thus, by careful "double-differencing" techniques against wire and air, the liquid's *relative* time of flight (ToF), phase, and RSSI can be measured with consistent accuracy. Since permittivity, refractive index, and loss factors are all defined relative to vacuum, such relative measurements align even better with the final requirements.

(2) The second opportunity arrives from the observation that phase can be derived from estimated ToFs, as well as from direct measurements. This serves as redundant information, which in turn can be utilized for coping with noise. Thus, *LiquidID* will use the modeled phase (from ToFs) to roughly estimate the integer ambiguity, then adjust the ambiguity from the measured phase, and finally refine the ToF from the precise ambiguity. This is a form of jointly estimating phase and ToF to arrive at the signal slowdown in liquid.

With these techniques, and subsequent stages of channel interpolation, container compensation, RSSI modeling, and others,

¹This is analogous to measuring the distance between point A and B with only the least significant digit of the measurement, say 6. We would not be able to tell if the actual distance was 6, 16, 26, 36, and so on. Measured phase is analogous to this least significant digit.

LiquidID arrives at an estimate of permittivity. The estimates are validated against specialized vector network analyzer (VNA) equipments (fortunately accessible in a Microwave Communications Lab).

Figure 1 shows our experimentation platform, with UWB antennas on two sides of the 2cm wide plastic container. The transmitter and receiver are not clock synchronized, neither do they need to be placed at any specific distance from the liquid. Regular multipath environments are acceptable, and any liquid can be poured. Our only requirement is that the liquid container be placed perpendicular to the wireless link, so that the signals are not incident obliquely onto the vertical cross-section of the liquid. This is important because oblique incidence triggers complex bending of signals at the boundary of the liquid, "smudging" the precision of our measurements. Of course, no special alignment is needed to ensure perpendicularity; just rough manual placement is more than adequate.




Figure 1: *LiquidID* experimental setup: (a) The UWB radios on two sides of the liquid-filled container placed perpendicularly. (b) The system schematic, with signals transmitted both through the antenna and wire.

We measure 33 different liquids, their permittivity ranging from the minimum $P_{air} = (1 + j0)$ to the maximum of $P_{water} = (75.7 + j14.3)$ (note that permittivity is a complex number). The 33 liquids cover the entire permittivity spectrum, with examples as: saline water, orange juice, pepsi, coke, isopropyl alcohol, ethanol, green tea, sweet tea, peanut oil, canola oil, etc. Median permittivity error is 9%, offering a resolution to separate sweet tea from green tea, orange juice from apple juice, olive oil from peanut oil, etc. We believe our permittivity errors are promising in comparison to specialized equipments (costing \$134,000) that quote 5% error using invasive probes dipped in the liquid. Our system, on the other hand, costs \$150, is non-invasive, and is practical for various applications.

In sum, this paper's contributions may be summarized as follows:

- *Identifying the viability of extracting material permittivity from the properties of signals passing through the material.* While we demonstrate our systems for liquids, the techniques generalize

to solids as well (with suitable modifications to the Physics models).

- *Developing synchronization-free techniques that quantify a signal's slow-down (inside a liquid) with high precision.* Building a robust system that does not require training or calibration, but identifies liquids with consistency.

We elaborate on these contributions next, starting with some background on permittivity, followed by overview, system design, and evaluation.

2 BACKGROUND ON PERMITTIVITY

The permittivity of a given liquid is technically a complex number ϵ^* as follows:

$$\epsilon^* = \epsilon' - j\epsilon'' \quad (1)$$

Here, ϵ' is the *dielectric constant* and ϵ'' is the *loss factor* of the liquid [37]. For vacuum, the loss factor is zero, therefore its permittivity is a real number denoted as $\epsilon_0 = \epsilon'_{vac}$. By convention, a liquid's permittivity is expressed relative to vacuum, as:

$$\epsilon^* = \left(\frac{\epsilon'}{\epsilon_0}\right) - j\left(\frac{\epsilon''}{\epsilon_0}\right) \quad (2)$$

Our final goal in this paper is to estimate these relative values denoted as $\epsilon'_r = \left(\frac{\epsilon'}{\epsilon_0}\right)$ and $\epsilon''_r = \left(\frac{\epsilon''}{\epsilon_0}\right)$. However, since neither of them can be measured directly, we have to rely on two other measurable quantities – *refractive index* and *attenuation* – to indirectly arrive at this estimation. To understand these indirect relationships, let us briefly look at refractive index and attenuation first.

Refractive Index (RI)

The refractive index n of a material is the ratio of the speed of light in vacuum, c , to the speed of electromagnetic waves in that material, v .

$$n = \frac{c}{v} \quad (3)$$

Even though a wave slows down in a given material, its frequency cannot change. As a result, the wave experiences a decrease in wavelength [9], dictated by:

$$f = \frac{c}{\lambda_0} = \frac{v}{\lambda} \quad (4)$$

where λ_0 is the wavelength of electromagnetic waves in vacuum and λ is the wavelength in that material. Now, refractive index is related to complex permittivity [37] as:

$$n = \frac{c}{v} = \frac{\lambda_0}{\lambda} = \sqrt{\frac{1}{2}\epsilon'_r \left\{ \sqrt{1 + \left(\frac{\epsilon''_r}{\epsilon'_r}\right)^2} + 1 \right\}} \quad (5)$$

This implies that refractive index alone is insufficient to estimate the 2 unknowns ϵ'_r and ϵ''_r . We need another equation to solve for complex permittivity.

Attenuation Factor (AF)

The “attenuation factor” of a material is defined as the width of the material needed to decay the strength of the electromagnetic field to $1/e = 0.368$ of its original (incident) value [19]. This width, α_d , is given by:

$$\alpha_d = \frac{\lambda_0}{2\pi} \sqrt{\frac{2}{\epsilon'_r \left(\sqrt{1 + \left(\frac{\epsilon''_r}{\epsilon'_r}\right)^2} - 1 \right)}} \quad (6)$$

For a liquid width of $d > \alpha_d$, the decay will be larger than $1/e$, and vice versa. Thus, the signal strength, SS , at the exiting boundary will be related to that at the entering boundary by:

$$\frac{SS_{exit}}{SS_{entry}} = \left(\frac{1}{e}\right)^{\frac{d}{\alpha_d}} \quad (7)$$

Measuring the LHS and knowing the d are adequate to obtain α_d . However, measuring the SS at precise liquid boundaries is not practical (since the antenna cannot be placed at the exact air-liquid interface). One potential solution is to measure the SS away from the boundary, and compute the reduction with and without liquid in the container. Unfortunately, this approach fails since signals also undergo reflections at liquid boundaries (as shown by gray arrows in Figure 2) and only a smaller fraction passes through the liquid. Unless we model these reflections, our estimates will be incorrect. We discuss this next and then arrive at an unified equation to obtain the final liquid permittivity.




Figure 2: Net signal attenuation after passing through liquid caused by (1) reflections at liquid boundaries, and (2) attenuation factor of the liquid.

Transmission Coefficient

When a traveling wave encounters a material boundary, only a portion of the incident energy penetrates into the new medium; the rest is reflected back. The fraction of penetrated energy is given by the transmission coefficient, t_E :

$$t_E = \frac{2Z_2}{Z_2 + Z_1} \quad (8)$$

where Z_1 is the intrinsic impedance of the material the waves are entering from, and Z_2 is the intrinsic impedance of the

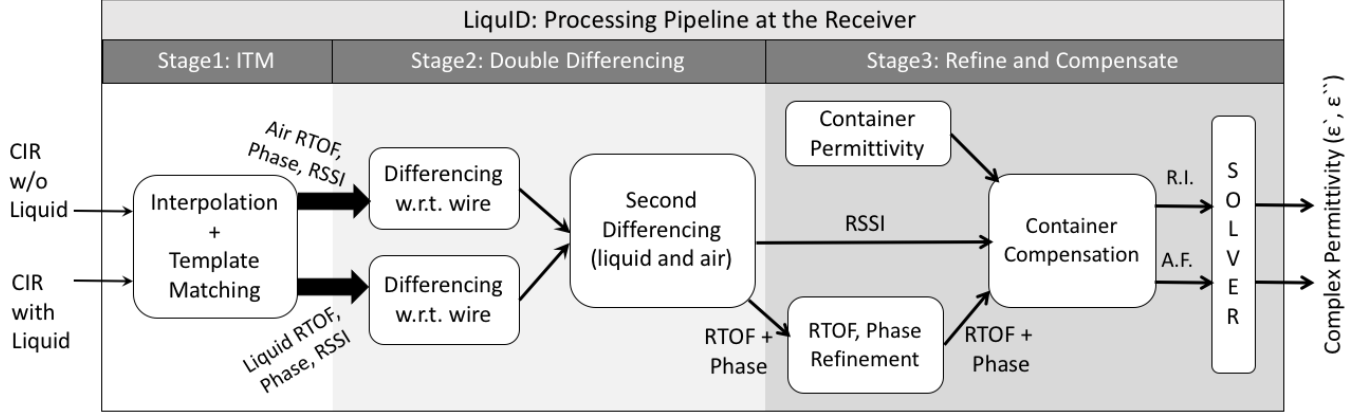


Figure 3: Four stages of processing, with Decawave's CIRs as inputs and liquid permittivity as output.

material the waves are entering into. Fortunately, intrinsic impedance is also related to permittivity as follows:

$$Z^* = \frac{Z_0}{\sqrt{\epsilon^*}} \quad (9)$$

where Z_0 is the impedance of free space, and ϵ^* is the permittivity of the entering material. Thus, it is now possible to combine all factors.

For this, we denote the transmission coefficient at the air-target boundary as $t_E^{air \rightarrow target}$ and that between the target-air boundary as $t_E^{target \rightarrow air}$. Then, we combine Equation 6 and Equation 8 to obtain the complete expression for attenuation of the signal before it reaches the receiver. The received signal strength (RSS) in the presence of the target liquid, RSS_{target} , compared to the RSS without the target, RSS_{air} is given by:

$$\frac{RSS_{target}}{RSS_{air}} = t_E^{air \rightarrow target} \cdot t_E^{target \rightarrow air} \cdot \left(\frac{1}{e}\right)^{\frac{d}{\sigma_d}} \quad (10)$$

Since Equation 10 also depends on both ϵ' and ϵ'' we solve it together with Equation 5 to obtain the complex permittivity of the target liquid, for a known liquid width of d . With this background, we are now ready to describe the LiquID system.

3 SYSTEM OVERVIEW

Figure 3 shows the computational pipeline underlying LiquID. We present an overview of the whole system first, followed by technical details of each stage.

Recall from Figure 1 that our system set-up is essentially a plastic container (to hold the liquid) with two UWB radios on either side. The radios are also connected with a wire to serve as a latency reference (explained later). The transmitter (Tx) forks the signal over the wire and the wireless antenna; the receiver (Rx) receives both over separate channels. To estimate permittivity, we perform 2 transmissions, first through a liquid-filled container, and then without the container (i.e., only air between the Tx and Rx). For each measurement, the Decawave receiver provides a discrete *channel impulse response* (CIR) shown in Figure 4.




Figure 4: A sample discrete CIR obtained directly from the Decawave UWB platform.

Given UWB's 1GHz wide bandwidth, the CIR taps (on the X-axis) are at 1 nanosecond time gaps. Nonetheless, Decawave performs fractional optimizations [12] and brings the ToF estimation down to ≈ 350 picoseconds. This implies that a distance computation using ToF would result in around 10cm error. While promising for applications like localization, separating liquids based on signal slowdown requires ≈ 30 picosecond resolution. This motivates the core problem of significantly improving ToF, in addition to carefully refining RSSI to finally achieve permittivity.

The LiquID receiver begins by accepting 2 CIRs as inputs – one through liquid and one through air. The processing on both the CIRs are identical, so let us consider only the CIR through liquid, denoted CIR_{liq} . Figure 4 shows an example CIR_{liq} – the first peak indicates the arrival of the wired signal while the second peak is due to the wireless signal (through the liquid).

In **Stage 1**, the goal is to identify the correct time at which the line of sight (LoS) signal arrives at the receiver. Given that Decawave's nanosecond resolution is inadequate, LiquID first performs a frequency-domain interpolation on CIR_{liq} . This inserts multiple samples between adjacent CIR taps, while also interpolating the phase of the samples. The next task is to identify the correct sample (in the interpolated CIR) corresponding to the LoS arrival. For this, LiquID extracts the portion of CIR corresponding to the wire – which captures the UWB hardware distortions – and correlates this CIR_{wire} against CIR_{liq} . The (sub-nanosecond) sample that spikes in correlation is denoted T_{liq} , and is declared as the time of LoS arrival through liquid.

The time of arrival for the wired signal is denoted T_{wire} , and was easier to detect since it was free of multipath.

Stage 2: To translate T_{liq} to absolute time of flight (ToF), the Rx must be precisely clock synchronized with the Tx. Although Decawave’s synchronization is sophisticated, it falls short of the needed 30 picosecond resolution. To completely sidestep synchronization, LiquID estimates *relative ToF* (RTOF) with respect to the wired path. The idea is that ToF through the wire remains very stable, hence $(T_{liq} - T_{wire})$ is essentially the relative ToF of the signal through the liquid. Similarly, $(T_{air} - T_{wire})$ is the relative ToF through air. Now, “differencing” again between the two quantities, (i.e., $(T_{liq} - T_{wire}) - (T_{air} - T_{wire})$) yields the relative ToF of liquid with respect to air (i.e., $(T_{liq} - T_{air})$). This, by design, is free of clock synchronization, and more importantly, captures the ToF *only* through the width of the liquid column (explained later).

Stage 3 focuses on refining ToF by leveraging signal phase. Observe that phase can be obtained from two sources: (1) direct measurement from the Rx, and (2) derived by dividing RTOF by λ , mod 2π . Of course, phase is affected by *integer ambiguity*, q (i.e., the signal may have traveled q full wavelengths during its time of flight). Thus, LiquID uses the (RTOF + derived phase) to obtain an estimate of q , refines q using the *measured* phase, and then refines relative ToF (RTOF) using the refined q . In a separate thread, LiquID also extracts the RSSI of the received signals, and again via “differencing”, computes the relative amplitude attenuation of liquid over air. However, to converge on actual permittivity, we still need to compensate the plastic container’s impact on RSSI.

Thus, as a one-time measurement, LiquID sends a signal through the empty container, computes the refined ToF, phase, and RSSI, and ultimately derives the permittivity of the container material. The container’s permittivity is fed into the processing pipeline so that the UWB signal’s behavior at the boundary of plastic and liquid can be modeled and compensated. The output from this “container compensation” module are the *refractive index* (RI) and the *attenuation factor* (AF) of the liquid, which are then fed into a MATLAB solver. This yields the final ϵ' and ϵ'' , and hence permittivity. We elaborate on the technical details of each stage, next.

4 SYSTEM DESIGN

Stage 1: Interpolation + Template Matching

A *channel impulse response* (CIR) describes the delays and strengths of various signal paths between Tx and the Rx. Figure 4 shows the CIR_{liq} from Decawave hardware – the resolution of the X-axis, a function of bandwidth, is 1 nanosecond [12]. The first task is to identify the time instant that corresponds to the arrival of the line of sight (LoS) signal path from the Tx to the Rx. Of course, this time may be somewhere between the available CIR samples (i.e., our CIR is only a measurement at the nanosecond granularity but the LoS signal could arrive at any time). This calls for a higher resolution CIR.

■ **Interpolation:** For higher resolution, we up-sample the CIR by adding *zeros* between existing samples and then apply a 1GHz low pass filter. Figure 5 shows this interpolated CIR. As evident from this example, the highest point in the interpolated CIR is actually shifted from the highest point of the un-interpolated CIR. In other words, interpolation takes the CIR closer to its original analog form, allowing for better estimation of the LoS delay and amplitude². This in turn allows for fractional (i.e., sub nanosecond) delay and more accurate amplitude estimation.

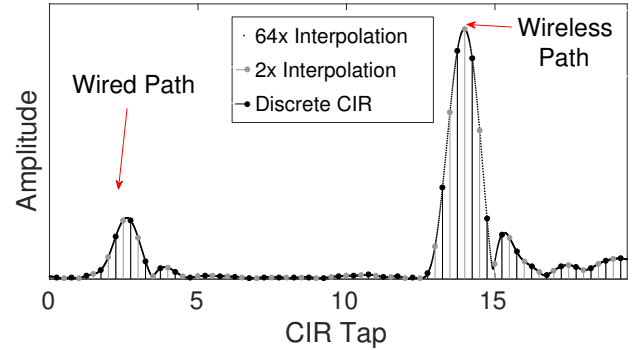


Figure 5: The interpolated version of a discrete CIR.

■ **Template Matching:** Of course, this estimation is still polluted by noise and various hardware distortions. Multipath, each adding a *Sinc(.)* function around its time of arrival may also add to the pollution. The CIR we have is a sampled version of these aggregated *Sinc()*s and hardware distortions. Thus, the peaks in our interpolated CIR may be spurious, i.e., the tallest sample may not correspond to the precise arrival time of the LoS signal.

To solve this problem, we utilize the CIR for the wired connection between the Tx and Rx. This CIR is free of multipath and captures the internal filtering effects (and other hardware distortions) of the UWB receiver. We call this the “template CIR” as shown in Figure 6 (and also subject it to interpolation). We then correlate the template CIR with the wireless CIR. The CIR sample at which the correlation spikes is declared, for now, as the *time of flight* (ToF) of the LoS signal.

In summary, we transmit UWB signals through a liquid-filled container, as well as through air, each yielding one CIR. Both CIR_{liq} and CIR_{air} are first interpolated followed by template matching. After matching against CIR_{liq} , we denote the best-match sample for the wired path as T_{wire} and the one for the wireless LoS path, as T_{liq} (see Figure 6). Similarly, for CIR_{air} , we get T_{wire} and T_{air} . Indeed, T_{liq} and T_{air} are a crude estimate of the ToF, however, we expect to improve this by making them relative.

Stage 2: Double Differencing

If the transmitter and receiver were synchronized, the antenna-to-antenna time of flight (ToF) could be calculated from

²Observe that interpolation also maintains the complex nature of the CIR, offering correct phase at the intermediate points after interpolation.

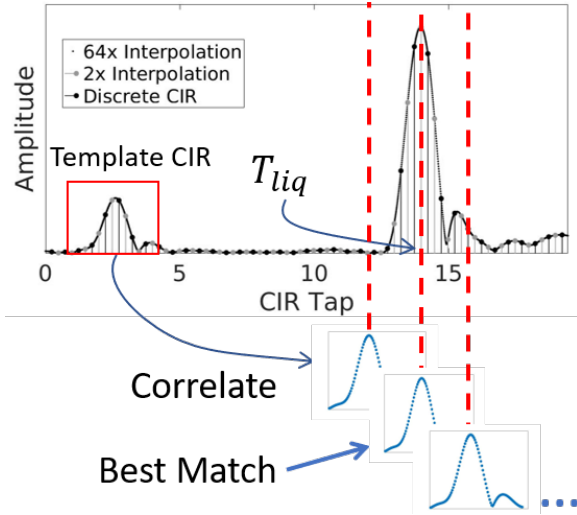


Figure 6: The CIR template from the wires channel used for correlating and finding the best sample, T_{liq} , at which LoS signal arrives through the liquid.

T_{liq} itself. However, achieving picoseconds level time synchronization is difficult even when using the same reference clock for two devices [8]. This motivates us to develop a synchronization-free technique. The simple opportunity is to utilize the wire delay (T_{wire}) as a fixed reference, and express the liquid or air path with respect to it. Specifically,

$$\begin{aligned}\Delta T_{wire}^{liq} &= T_{liq} - T_{wire} \\ \Delta T_{wire}^{air} &= T_{air} - T_{wire}\end{aligned}\quad (11)$$

While this obviates the need for synchronization, it is still inadequate for estimating ToF through liquid. The reason is that ΔT_{wire}^{liq} is not just the signal propagation through the liquid column, but also includes propagation delays through antenna connectors, air gaps between antennas and the liquid, and even through the material of the container. Figure 7 illustrates the path and we model this as:

$$\Delta T_{wire}^{liq} = \frac{2L_{ant}}{v_{wire}} + \frac{L_{air}^{(1)}}{c} + \frac{L_{liq}}{v} + \frac{L_{air}^{(2)}}{c} - \frac{L_{wire}}{v_{wire}} \quad (12)$$

where v is signal velocity in the liquid, L_{liq} is the liquid's column depth, and $L_{air}^{(1)}$ and $L_{air}^{(2)}$ are the first and second air gaps before and after the liquid column.

Similarly, for ΔT_{wire}^{air} , shown in Figure 8, we can write a similar equation as:

$$\Delta T_{wire}^{air} = \frac{2L_{ant}}{v_{wire}} + \frac{L_{air}}{c} - \frac{L_{wire}}{v_{wire}} \quad (13)$$

where L_{ant} is the length of each antenna connector, L_{air} is the physical air gap between antennas, L_{wire} is length of the wired path, v_{wire} is speed of the signals in wire, and c is the speed of UWB signals in air.

To estimate the ToF *only* within the liquid column, we need to isolate out the term (L_{liq}/v) from Equation 12. To achieve this, we perform a second differencing operation between the




Figure 7: The ToF consists of various delay components in addition to delay through liquid alone.




Figure 8: The ToF consists of other delay components in addition to delay through air alone.

above two equations as follows.

$$\begin{aligned}\Delta T_{air}^{liq} &= \Delta T_{wire}^{liq} - \Delta T_{wire}^{air} \\ \Delta T_{air}^{liq} &= \frac{L_{air}^{(1)}}{c} + \frac{L_{air}^{(2)}}{c} + \frac{L_{liq}}{v} - \frac{L_{air}}{c}\end{aligned}\quad (14)$$

Since $(\frac{L_{air}^{(1)}}{c} + \frac{L_{air}^{(2)}}{c} - \frac{L_{air}}{c}) = (-\frac{L_{liq}}{c})$, we have the ToF through the liquid, relative to that of air, as:

$$Relative\ ToF = \Delta T_{air}^{liq} = \frac{L_{liq}}{v} - \frac{L_{liq}}{c} \quad (15)$$

During an actual experiment, many values of ΔT_{wire}^{air} are obtained first (for about a minute) and then the liquid column is introduced. The median ΔT_{wire}^{air} is used for the differencing in Equation 14. Since the liquid column depth (L_{liq}) is accurately known, we solve Equation 15 using experimentally obtained values of ΔT_{air}^{liq} and estimate the velocity v of signals in the liquid. This v is essentially the slower signal velocity in liquid that we have been seeking to extract. We will improve this accuracy in the next stage and then plug the values in the LHS of Equation 5 to obtain refractive index.

Double Differencing Phase

Observe that the differencing operations hardly introduce errors in relative ToF; the root of ToF errors are still sourced in the original interpolation and template matching operations (i.e., in the estimation of T_{liq} and T_{air}). To refine this in stage 3 of the processing pipeline, we will use the phase corresponding to the samples T_{liq} and T_{air} , denoted ϕ_{liq} and ϕ_{air} , respectively. Phase brings value because it is an attribute of the original infinite bandwidth impulse that was obtained even

before filtering at the receiver, hence, the phase under the CIR peaks are relatively constant as seen from Figure 9. The phase variations over the peak are about 1.5° in the wired path and around 7° in the air path. Therefore, inaccuracies in locating the exact peak do not significantly affect the phase, but can dramatically alter the ToF estimate. This is the key opportunity. However, to apply phase to relative ToF, we require similar double differencing operations to be imposed on phase as well.




Figure 9: Phase is stable for many adjoining locations at the prominent peaks. An error in the estimation of the correct CIR tap results in a small phase error, but large ToF error.

Differencing phase is almost identical to ToF and can be expressed as follows.

$$\begin{aligned}\phi_{wire}^{liq} &= \phi_{liq} - \phi_{wire} \\ \phi_{wire}^{air} &= \phi_{air} - \phi_{wire}\end{aligned}\quad (16)$$

Once computed relative to the wire, double differencing over the relative phases ϕ_{wire}^{air} and ϕ_{wire}^{liq} will provide the phase difference between the liquid and the air peaks:

$$\Delta\phi_{air}^{liq} = \phi_{wire}^{liq} - \phi_{wire}^{air}\quad (17)$$

Of course, the phase may wrap (called *integer ambiguity*), but we use the *relative ToF* estimate to unwrap it in the next stage.

■ Double Differencing RSSI

Permittivity is derived from *refractive index* (RI) and *attenuation factor* (AF). Recall that RI is essentially the slowdown of a signal in a medium, which necessitates relative ToF. However, to estimate AF, LiquiD also needs to measure the signal's amplitude degradation in the liquid, w.r.t. air. Importantly, the absolute values of amplitude vary substantially due to small variations in the AGC gains at the receiver. Figure 10 shows this variation in the absolute received signal strength (RSS) for different packets. To handle this, we require a differencing treatment on RSSI as well, similar to ToF.

The reference wired path is again useful for comparison. For every packet, the amplitude of the T_{liq} sample is normalized with the amplitude of T_{wire} . Similarly amplitude for T_{air} is normalized with the corresponding T_{wire} amplitude. In essence, we scale all CIRs with the amplitude of their respective T_{wire} . The relative received signal strengths (relative RSS) can now




Figure 10: Large variation of absolute RSS, however, variation much less for relative RSS.

be modeled as:

$$\begin{aligned}RSS_{wire}^{liq} &= \frac{RSS_{liq}}{RSS_{wire}} \\ RSS_{wire}^{air} &= \frac{RSS_{air}}{RSS_{wire}}\end{aligned}\quad (18)$$

Figure 10 shows how the variation in the relative RSS (across many packets) is much smaller now, while Figure 11 shows very good alignment of 100 (air) CIRs after scaling and aligning at sample T_{wire} .

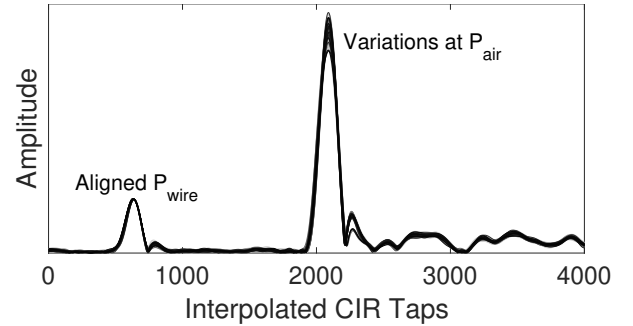


Figure 11: 100 overlapped “air” CIRs after scaling by the amplitude of T_{wire} . This demonstrates the opportunity to accurately compare liquid with air.

Finally, to arrive at the attenuation factor (AF), we difference a second time between the liquid's and air's RSS, giving us:

$$RSS_{air}^{liq} = \frac{RSS_{liq}}{RSS_{air}}\quad (19)$$

Figure 12 shows the result of this stage which gives the relative ToF and the relative RSS. Observe that T_{wire} serves as a precise reference against which the liquid and air can be compared, both in terms of ToF and signal strength. We can now replace the LHS of Equation 7 (i.e., the attenuation factor equation) with actual RSSI values obtained from our experiments.

Stage 3: ToF Refinement and Container Compensation

LiquiD is now ready to fine tune the relative ToF (RTOF) using the relative phases, and also compensate for the effect of the plastic container on the measurements.




Figure 12: The wired reference aligns T_{wire} in the CIR, thereby enabling relative ToF and relative RSS between liquid and air.

Relative-ToF Refinement

Since phase exhibits better precision than relative ToF (discussed earlier for Figure 9), we see an opportunity for refining RTOF. Observe that the RTOF is a function of the distance traveled by the signal in the liquid, which in turn equals an integer multiple of λ + measured phase $\hat{\phi}$. Concretely:

$$RTOF = \Delta T_{air}^{liq} = (q + \frac{\hat{\phi}(\text{mod}2\pi)}{2\pi}) \frac{\lambda_0}{c} \quad (20)$$

where q is an integer number of wavelengths wraps, and $\hat{\phi}$ is calculated from the RTOF (i.e., $\hat{\phi} = \frac{RTOF \cdot c}{\lambda_0}$). Also, λ_0 is the wavelength of the UWB signal in air, and c is the speed of UWB signals in air. A unique solution to this equation can be obtained using the constraints that q is an integer and $\hat{\phi}$ is bounded in $[0, 2\pi]$.

Now, we compare $\hat{\phi}$ with $\Delta\phi_{air}^{liq}$ that we had computed in stage 2 via double differencing. These should be close, but if not, we adjust q up or down to bring them closer. For instance, if $\hat{\phi} = 10^\circ$ and $\Delta\phi_{air}^{liq} = 350^\circ$, then we reduce q by 1. In other words, we trust $\Delta\phi_{air}^{liq}$ more and believe that the actual signal path must be $(q-1)\lambda_0 + 350^\circ$, as opposed to $q\lambda_0 + 350^\circ$. Thus, with this adjusted q , denoted \hat{q} , we refine RTOF as:

$$RTOF_{refined} = (\hat{q} + \frac{\Delta\phi_{air}^{liq}}{2\pi}) \frac{\lambda_0}{c} \quad (21)$$

This refined RTOF is converted to velocity of UWB signals in the liquid, and averaged over 20 samples to suppress noise. As a final step, we obtain the estimated Refractive Index from Equation 3 as follows:

$$n = \frac{c}{v} = \frac{(\hat{q} + \frac{\Delta\phi_{air}^{liq}}{2\pi})\lambda_0}{L_{liq}} + 1 \quad (22)$$

Container Compensation

Of course, the container used for storing the liquid also influences the ToF, phase, and the RSS. Fortunately, we can perform all of the above measurements on an empty plastic container and obtain its complex permittivity as well. We observed a phase difference of around 45° and a negligible attenuation factor by introducing the empty plastic container. This results in the complex permittivity of $(3.45 - 0j)$ that closely matches typical material properties for acrylic [15]. We subtract the

equivalent plastic ToF from all our measurements and modify Equation 10 as follows:

$$\frac{RSSI_{liq}}{RSSI_{air}} = t_E^{air \rightarrow c} \cdot t_E^{c \rightarrow liq} \cdot t_E^{liq \rightarrow c} \cdot t_E^{c \rightarrow air} \cdot \left(\frac{1}{e}\right)^{\frac{L_{liq}}{\alpha_d}} \quad (23)$$

With these compensations performed, we now have the values of the RSSI ratio and the expected Refractive Index of the liquid. For a different container, its permittivity, thickness, and liquid width must be taken into account and the wireless signals must be normally incident on the face of the container.

Solve for Complex Permittivity

To solve for permittivity, we rely on Equations 5 and 10). Except ϵ' and ϵ'' , we know all the parameters in these equations (note that $\lambda = \lambda_0/n$, and we know the refractive index n at this point). We can therefore solve for ϵ' and ϵ'' – we employ the Matlab solver (lsqnonlin) and arrive at the best estimates. The next section reports on the accuracy of our estimates.

EVALUATION

5.1 Experimental Setup

Our experimental setup for evaluating LiquID is shown in Figure 1. We use the Decawave UWB Trek1000 [18] evaluation kit for our experiments. The cost of the hardware is \approx \$250. However, the cost of the UWB chip is under \$10 [3] and can be integrated to use a mobile device CPU as host controller [6]. The liquid is placed inside an acrylic container located between the Tx and Rx UWB radios. We place the container's surface perpendicular to the wireless link to ensure that the wireless signal does not undergo bending and diffraction.³ Note, that this does not require calibration and manual placement is sufficiently good to avoid diffraction.

The container we use in our experiments has a depth of 2 cm, breadth of 38 cm, and a height of 36 cm. In general, the depth should be sufficiently large to create enough measurable slow down in the signal without significantly weakening the RSSI of the signal. The length and breadth on the other hand need to be larger than the signal wavelength to avoid signal diffraction.

Our transmitter (Tx) sends a known UWB packet every 500ms at 4GHz center frequency. A receiver (Rx) records the channel impulse response (CIR) from the received packet preamble, and exports it along with the ToF, phase, and RSSI. The duration between packets can be reduced to about 30ms allowing just enough time to collect the interesting portion of the CIR off the UWB receiver. The output of the receiver is collected on a Lenovo laptop using a USB serial cable where it is processed using LiquID's three stages in MATLAB. The Tx and Rx are also connected directly with a wire—the CIR collected includes this wired link as well.

All experiments are run in an office setting with standard furniture and multipath effects. All liquids including milk were at room temperature between 23°C and 26°C . Effervescent

³In cases where the container is circular, the wireless link should be perpendicular to the tangent at the incident point.

liquids were allowed to stand several hours in an open mouth vessel to remove all bubbles before testing. Most of the liquids were directly procured from popular grocery stores. NaCl and Glucose solutions of different concentrations, however, were prepared by mixing non-iodized salt and granulated sugar in distilled water in our lab.

Baseline: As baseline for estimating permittivity, we use a vector network analyzer. The experimental setup is shown in Figure 13. The setup uses an Agilent N5242A PNA-X Microwave Network Analyzer [7] which costs over \$134,000. The PNA is connected to an HP-85070A dielectric probe [5]. The probe is dipped into the liquid which enables accurate measurement of S11 parameters. A proprietary software uses these parameters to compute the permittivity of the liquid sample. The permittivity measurements from the network analyzer are known to have a median error of 5% [33].




Figure 13: Baseline: Network Analyzer Experimental Setup

5.2 Evaluation Metrics

We evaluate LiquiD along two metrics:

Permittivity Estimation: We measure the complex permittivity, $\epsilon^* = \epsilon' - j\epsilon''$, using LiquiD for 33 different liquids shown in Table 1. We compare with the VNA baseline and report the mismatch between the two.

Liquid Identification: We use our permittivity estimates from LiquiD to identify and distinguish between the 33 liquids. We compute and report the confusion matrix for identifying all these 33 liquids.

In addition to the above metrics, we report several micro-benchmark results such as ToF measurement accuracy, refractive index and attenuation factor coverage across all kinds of liquids. These micro-benchmarks give us additional insights into LiquiD's performance.

5.3 Results

5.3.1 Estimating the Permittivity. We report the permittivity estimates for 33 liquids and compare our results with the baseline network analyzer measurements. Table 1 shows both

LiquiD's and the baseline's estimates for both ϵ' and ϵ'' . Figures 14 and 15 show the mismatch between LiquiD and the baseline. The median mismatch in dielectric constant ϵ' is 3.29. For ϵ'' the median mismatch is 0.99. This translates into a median relative error of 9% in ϵ' and 11.9% in ϵ'' , which are reasonably comparable to the baseline's error of 5%. While LiquiD has slightly higher error, the LiquiD setup is four orders of magnitude cheaper than the baseline setup and is non-invasive i.e. it does not require placing a probe inside the liquid.

Liquid	LiquiD		Baseline	
	ϵ'	ϵ''	ϵ'	ϵ''
Distilled Water	72.34	13.16	75.64	14.7
Mineral Water	75.79	14.87	74.79	14.83
Orange Juice	55.28	18	67.38	18.77
Apple Juice	67.33	15	74.01	14.29
Grape Juice	67.06	20.43	67.16	18.74
Pepsi	71.92	17.24	74.92	20.98
Coca-Cola	67.1	16.93	71.25	18.87
Diet Pepsi	76.3	14.41	75.1	14.85
Diet Coca-Cola	79.09	14.16	80.79	19.39
NaCl 0.1M	67.75	18.63	77.68	23.96
NaCl 0.05M	65.07	15.99	75.16	18.14
Glucose 5%	54.57	14.72	71.52	15.70
Glucose 10%	65.52	15.76	68.65	16.19
Skim Milk	73.96	19.02	68.09	17.41
Milk 1%	76.22	19.34	66.87	17.12
Milk 2%	73.94	19.22	66.53	17.05
Whole Milk	70.31	19.38	64.34	16.51
Sweet Tea	77.04	16.26	71.23	16.72
Green Tea	75.26	15.65	71.54	16.28
Cranberry Juice	72.36	17.35	69.39	17.45
Coffee	68.42	12.71	73.49	12.97
Pine Sol	70.23	16.41	72.52	18.55
Isopropanol 50%	26.95	15.26	30.4	19.79
Isopropanol 70%	8.22	10.09	16.21	12.87
Isopropanol 91%	5.37	4	5.35	4.21
Ethanol 70%	15.95	12.84	17.63	15.11
Vinegar	13.88	5.46	26.84	2.17
Peanut Oil	2.22	0	2.6	0.12
Olive Oil	2.25	0	2.87	0.13
Cutting Oil	1.99	0	1.35	0
Soyabean Oil	2.25	0	2.63	0.12
Motor Oil	1.92	0	2	0
Corn Oil	2.28	0	2.57	0.14

Table 1: Measured permittivity with LiquiD vs. Network Analyzer for various liquids.

Table 2 compares the complex permittivity measured by LiquiD with that available in literature for some liquids [13, 21, 26, 27]. The table shows that LiquiD reports permittivity measurements very close to the expected values.

In order to better understand where the error comes from, we plot the refractive index and attenuation factor of the various liquids as shown in Figures 16 and 17. Recall that the permittivity coefficients ϵ' and ϵ'' are estimated from the refractive

Liquid	LiquID		Literature	
	ϵ'	ϵ''	ϵ'	ϵ''
Distilled Water	72.3	13.1	75.7	14.3
Ethanol 70%	15.95	12.84	16	14
Isopropanol 91%	5.37	4	5.5	4.8
Soybean Oil	2.25	0	2.71	0.174

Table 2: Measured permittivity with LiquID vs. that reported in literature.




Figure 14: Mismatch between LiquID and baseline in ϵ'




Figure 15: Mismatch between LiquID and baseline in ϵ''

index n and attenuation distance α_d using equations 5, 7, and 10. However, these equations are non-linear and hence any small error in measuring the refractive index or attenuation loss can lead to a large error in estimating the permittivity.

The figures also show that the liquids can be clustered into three types (Ordering of liquids in Figure 16 has been sorted by the refractive index to make this point clear):

- **Oils:** have a small refractive index and permittivity but a large attenuation distance and hence do not slow down or attenuate the signal much. They are *transparent* to radio frequencies.
- **Alcohols:** have a medium refractive index, attenuation distance and permittivity.
- **Water based liquids:** have a large refractive index and permittivity and small attenuation distance. They are close to *opaque* to radio frequencies.

Figure 18 also shows the spread of the three types of liquids over the permittivity spectrum. The figure shows that we cover a large range of permittivity values along both ϵ' and ϵ'' .

5.3.2 Liquid Identification. Figure 19 shows the confusion matrix for identifying and distinguishing the 33 liquids using permittivity estimates from LiquID. Each liquid was tested 10




Figure 16: Refractive index of tested liquids




Figure 17: Attenuation factor of tested liquids




Figure 18: The Oils, Alcohols, and Water solutions span a large spectrum of permittivity values.

times and compared with the estimates from one initial experiment. Evidently, LiquID identifies most of the liquids correctly. While confusion does occur, it is limited to liquids with very close permittivity like mineral water and diet peps, where the discrepancy in the actual permittivity values is within the margin of measurement error. Even if we use the specialized VNA system, confusion between such liquids will persist since their permittivity values are extremely close (i.e., within the VNA's 5% error range).

Figure 20 also shows the confusion matrix for identifying the same liquids by directly using RSSI and ToF (as opposed to permittivity). Evidently, the identification performs worse, exhibiting much more confusion between the liquids. This emphasizes the importance of using permittivity as the physical attribute for liquid identification. While the permittivity is derived from RSSI and ToF, recall from equations 5 and 6 in section 2, that the relation between permittivity and RSSI/ToF is a complex

non-linear equation. Clearly, this non-linearity is important in developing the signature necessary for liquid identification.




Figure 19: Confusion matrix for liquid identification with LiquiD, using permittivity.

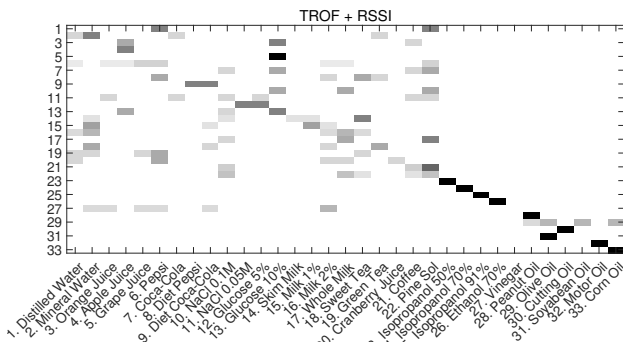


Figure 20: Confusion Matrix for liquid identification based only on RSSI and ToF measurements.

5.3.3 Time-of-flight Micro-benchmarks. In this section, we micro-benchmark LiquiD’s time of flight (ToF) estimation process. First, we evaluate the precision of LiquiD’s ToF estimation using RTOF described in section 3. We compare it to two baselines. The first is provided directly by the Decawave UWB kit. The kit is designed to enable ranging between two devices and hence it provides ToF measurements for the first arriving path at the receiver antenna. The second baseline is a standard max-peak finder that picks the first maximum peak as the first path. Figure 21 shows a CDF of the precision in estimating the time of flight for each of the compared schemes. LiquiD’s RTOF achieves a median precision of 55ps that is 3x better than Decawave and 2x better than max-peak finder.

We now evaluate the importance of LiquiD’s 3 stage system design in improving the estimation of the ToF. We compare results from the output of each stage for LiquiD. The first is LiquiD using RTOF. The Second is LiquiD using RTOF and interpolation. Finally, the third is LiquiD using RTOF, interpolation and phase based refinement i.e., the entire LiquiD system. Figure 22 shows the precision error in the ToF for the 3 schemes. Using all three stages, the median error is less than 1.7ps which is 31x lower than simply using RTOF. Even after highly upsampling the CIR using interpolation, the precision is still around




Figure 21: Comparison between Decawave’s ToF value precision and LiquiD’s ToF estimation precision




Figure 22: Comparison between precision obtained for ToF at various stages of LiquiD’s processing.

81.49ps at the 90%ile. This benefit translates to a significantly improved estimation ability for complex permittivity.

6 RELATED WORK

Optical Spectroscopy: Optical spectroscopy entails analyzing the absorption and emission properties of materials in light frequencies [23, 24]. Although very precise, the specialized equipments are bulky and expensive (\$30,000+), making them unsuitable for low cost operation in a ubiquitous setting. An interesting approach to liquid identification is proposed in [34] that uses photoacoustic effect. Their approach relies on shining lights of various wavelengths through liquids. However, penetration of light in dark colored liquids is limited. Instead, our approach uses radio frequency waves that can penetrate dark colored liquids including thick black oils. Moreover, [34] shows promise in a very limited set of water based solutions. In contrast, this work identifies liquids across various classes such as water solutions, oils, and alcohols. In its core, LiquiD measures a fundamental physical property of all substances and does not just classify liquids.

Impedance spectroscopy at RF: These techniques are based on impedance responses of a material to an applied RF field (example techniques include co-axial probe methods, free space methods, transmission line and reflection methods, resonant techniques [1, 11, 14, 17, 41]). A popular approach uses a co-axial probe dipped into the liquid while a vector network analyzer (VNA), connected to the probe, measures permittivity [33]. Impedance mismatch at the probe-material boundary causes reflections whose properties determine the liquid’s permittivity. While useful for studying permittivity across a broad RF spectrum, the equipment is not only expensive and bulky,

but also invasive. In contrast LiquID provides a non-invasive and low cost alternative.

UWB and RFID Approaches: Prior works have used RFID and UWB wireless signals to identify material properties such as oil adulteration[28], attenuation due to building materials[25, 38] (wood, gypsum, glass), and salt concentration[35]. Closest to our work is Tagscan [39], which uses RFID signals to classify 10 liquids. Tagscan’s approach can only classify liquids whose phases differ by certain relationships. Specifically, two liquids that cause different number of 2π phase wraps to the signal, cannot be discriminated. Thus, the classical problem of *integer ambiguity* still remains, allowing TagScan to discriminate a few selected liquids. In contrast, LiquID exploits high resolution ToF of UWB to resolve the integer ambiguities, and in combination with precise double-differencing methods, estimates the actual complex permittivity ($\epsilon' + \epsilon''$). The technique scales to any liquid (33 liquids reported in this paper), while requiring no training for classification – the estimated permittivity can be directly compared against ground truth.

ToF, Phase, and RSSI: A number of localization, and gesture tracking papers leverage from ToF, phase, and RSSI measurements not only in RF [10, 16, 29, 31, 40], but also in acoustics [30, 32, 42]. SAIL [31] uses ToF measurements for indoor localization. WiTrack [10] detects human motion based on ToF computed from FMCW Radar[10]. SafetyNet [20], and [40] uses phase for computing orientation of drones and objects. In contrast, the ToF, phase, and RSSI accuracy requirements for LiquID is very high, in the granularities of few *mm* and 0.3 dB. LiquID uses a multi stage signal processing pipeline to achieve its requirements.

7 POINTS OF DISCUSSION

■ Closing the Gap in Permittivity Error.

Permittivity error is sourced in ToF, phase, and RSSI errors. While improvements may be possible via higher sampling frequency (i.e. bandwidth) or longer point FFTs, we focus on the room for algorithmic improvement. Specifically, ToF and phase errors can be improved by further optimizing the choice of the CIR sample at which the signal arrived. Of course, this is a function of multipath and diffraction. While we partially compensate for multipath, diffraction is harder since diffracted paths could arrive at the same time (or even earlier) than the slowed-down path through the liquid. This also affects RSSI errors, especially because the path through the liquid is already weak. Isolating diffraction is hard; we leave this to future work.

■ Can Any Sized Container be Used?

The dimensions of our container are chosen carefully so that signals do not impinge obliquely on the vertical surface, while also ensuring that RSSI and slow-down are reasonable. Arbitrary shaped containers are extremely challenging (according to [37], “the general case of arbitrary dielectrics and oblique incidence is of great complexity”). Moreover, when the container is below a threshold size, the diffracted signals may completely

drown the through-liquid signal. LiquID for arbitrary containers remains an open problem.

■ Permittivity versus Frequency.

Signal slowdown, and hence permittivity, is a function of frequency. For our high bandwidth signals, the measured permittivity is actually an aggregate of many frequency-specific permittivities. With narrow band signals, the permittivity estimates can be improved at the expense of loosing ToF precision (recall high bandwidth is necessary to precisely estimate ToF). This is a tradeoff and perhaps new techniques are feasible that leverage the best of both worlds. We leave this to future work.

■ Beyond Liquids.

Our core techniques should generalize to solids, with suitable modifications to the physics model and parameters. In fact, the signal slowdown may be more pronounced, aiding in more precise ToF calculation. The only technical issue may be around RSSI – solids may attenuate RSSI more than liquids, requiring a higher-power signal, or narrower material width. A follow up work needs to thoroughly address the case for solids.

■ Considerations for a Mobile Form Factor.

Fitting LiquID into a mobile form factor entails incorporating the space, energy, and computation constraints of such platforms. An UWB chip occupies a $4mm \times 4mm$ area and weighs 0.105gm [4]. Whereas we have used half-wavelength dipole antennas in our experiments, much smaller antennas are available [2]. UWB is a low-power protocol [22] and the decawave chip is rated to consume a maximum of 120mAh [4]. The signal processing blocks required by LiquID are already present on mobile devices [36]. Pipelining the signal processing with fetching of the CIR data can allow LiquID to run at near-realtime. Finally, we envision a mobile device with an antenna connected to an extensible wire. A liquid container is placed between the mobile device and this antenna to identify the liquid.

8 CONCLUSION

This paper shows the feasibility of identifying liquids by analyzing UWB signals passing through it. We measure the time of flight of the signal and combine with its phase and RSSI to ultimately model the permittivity of the liquid. Given permittivity serves as a signature, it is now possible to identify the liquid without inserting probes into it. Our solution is realtime (sub-second latency), cheap (\approx \$150), and lightweight (few pounds), underpinning a variety of applications. Our next step is focussed on analyzing more complex liquid mixtures, such as impure drinking water, blood, saliva, and other biologically relevant liquids.

ACKNOWLEDGMENTS

We would like to acknowledge Prof. José Schutt-Ainé for enabling the use of a vector network analyzer and dielectric probe required to obtain the baseline permittivity. We would also like to thank Davy Davidson of Boda Acrylic for creating the experimental acrylic box at short notice. We sincerely thank our shepherd, Dr. Tam Vu, and the anonymous reviewers for their valuable feedback. We are grateful to NSF (CNS - 1719337) for partially funding this research.

REFERENCES

- [1] 6 techniques for measuring dielectric properties. <https://www.degruyter.com/downloadpdf/books/9783110455403/9783110455403-007/9783110455403-007.pdf>.
- [2] AccuraUWB Flex Series. <http://cdn.taoglas.com/datasheets/FXUWB20.01.0100C.pdf>.
- [3] DW1000-I-TR13 individual uwb chip pricing. <https://www.digikey.com/short/j47cd1>.
- [4] DW1000 IEEE802.15.4-2011 UWB Transceiver Datasheet. <https://www.decawave.com/products/dw1000>.
- [5] HP85070A Dielectric Probe Kit. <https://literature.cdn.keysight.com/litweb/pdf/85070-90001.pdf?id=1326230>.
- [6] Moving from trek1000 to a product. https://www.decawave.com/sites/default/files/aps016_moving_from_trek1000_to_a_product.pdf.
- [7] N5242a PNA-X Microwave Network Analyzer. <https://literature.cdn.keysight.com/litweb/pdf/N5242-90007.pdf?id=1118335>.
- [8] National instruments ni-tcl technology for timing and synchronization of modular instruments. <http://www.ni.com/tutorial/3675/en/>.
- [9] Reflection and refraction. http://electron6.phys.utk.edu/optics421/modules/m1/reflection_and_refraction.htm.
- [10] ADIB, F., KABELAC, Z., KATABI, D., AND MILLER, R. C. 3d tracking via body radio reflections. In *NSDI* (2014), vol. 14, pp. 317–329.
- [11] AFSAR, M. N., BIRCH, J. R., CLARKE, R., AND CHANTRY, G. The measurement of the properties of materials. *Proceedings of the IEEE* 74, 1 (1986), 183–199.
- [12] ALARIFI, A., AL-SALMAN, A., ALSALEH, M., ALNAFESSAH, A., AL-HADHRAMI, S., AL-AMMAR, M. A., AND AL-KHALIFA, H. S. Ultra wideband indoor positioning technologies: Analysis and recent advances. *Sensors* 16, 5 (2016), 707.
- [13] ANDERSON, C. *Determining the complex permittivity of materials with the Waveguide-Cutoff method*. PhD thesis, 2006.
- [14] BAKER-JARVIS, J., VANZURA, E. J., AND KISSICK, W. A. Improved technique for determining complex permittivity with the transmission/reflection method. *IEEE Transactions on microwave theory and techniques* 38, 8 (1990), 1096–1103.
- [15] BLATTENBERGER, K. Dielectric Constant, Strength, and Loss Tangent.
- [16] CAMPBELL, B., DUTTA, P., KEMPKE, B., KUO, Y.-S., AND PANNUTO, P. Decawave: Exploring state of the art commercial localization. *Ann Arbor 1001* (2015), 48109.
- [17] COURTNEY, C. C. Time-domain measurement of the electromagnetic properties of materials. *IEEE Transactions on Microwave Theory and Techniques* 46, 5 (1998), 517–522.
- [18] DECAWAVE. Decawave. <http://www.decawave.com/>.
- [19] FEYNMAN, R. P., LEIGHTON, R. B., AND SANDS, M. *The Feynman lectures on physics, Vol. I: The new millennium edition: mainly mechanics, radiation, and heat*, vol. 1. Basic books, 2011.
- [20] GOWDA, M., MANWEILER, J., DHEKNE, A., CHOUDHURY, R. R., AND WEISZ, J. D. Tracking drone orientation with multiple gps receivers. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking* (2016), ACM, pp. 280–293.
- [21] GREGORY, A. P., AND CLARKE, R. *Tables of the complex permittivity of dielectric reference liquids at frequencies up to 5 GHz*. National Physical Laboratory Teddington, 2001.
- [22] GUTIERREZ, J. A., CALLAWAY, E. H., AND BARRETT, R. *IEEE 802.15.4 Low-Rate Wireless Personal Area Networks: Enabling Wireless Sensor Networks*. IEEE Standards Office, New York, NY, USA, 2003.
- [23] HIND, A. *Agilent 101: An introduction to optical spectroscopy*, 2013.
- [24] JAMES, J. F., STERNBERG, R. S., AND RICE, S. A. The design of optical spectrometers. *Physics Today* 23 (1970), 55.
- [25] JATURATUSSANAI, P., CHAMCHOY, M., AND PROMWONG, S. Characteristics of uwb propagation through building materials. In *Communications and Information Technology, 2005. ISCT 2005. IEEE International Symposium on* (2005), vol. 2, IEEE, pp. 987–990.
- [26] KAATZE, U. Complex permittivity of water as a function of frequency and temperature. *Journal of Chemical and Engineering Data* 34, 4 (1989), 371–374.
- [27] KOMAROV, V., WANG, S., AND TANG, J. Permittivity and measurements. *Encyclopedia of RF and microwave engineering* (2005).
- [28] LEVITAS, B., MATUZAS, J., VISWANATH, G., BASALINGAPPA, V., AND VENKOPARAO, V. Uwb based oil quality detection. In *Ultra-Wideband (ICUWB), 2011 IEEE International Conference on* (2011), IEEE, pp. 220–224.
- [29] MA, Y., HUI, X., AND KAN, E. C. 3d real-time indoor localization via broadband nonlinear backscatter in passive devices with centimeter precision. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking* (2016), ACM, pp. 216–229.
- [30] MAO, W., HE, J., AND QIU, L. Cat: high-precision acoustic motion tracking. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking* (2016), ACM, pp. 69–81.
- [31] MARIAKAKIS, A. T., SEN, S., LEE, J., AND KIM, K.-H. Sail: Single access point-based indoor localization. In *Proceedings of the 12th annual international conference on Mobile systems, applications, and services* (2014), ACM, pp. 315–328.
- [32] NANDAKUMAR, R., IYER, V., TAN, D., AND GOLLAKOTA, S. Fingerio: Using active sonar for fine-grained finger tracking. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016), ACM, pp. 1515–1525.
- [33] NOTE, A. A. *Agilent basics of measuring the dielectric properties of materials*. *Agilent Literature Number* (2006).
- [34] RAHMAN, T., ADAMS, A. T., SCHEIN, P., JAIN, A., ERICKSON, D., AND CHOUDHURY, T. Nutrilizer: A mobile system for characterizing liquid food with photoacoustic effect. In *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM* (2016), ACM, pp. 123–136.
- [35] SELMIC, R. R., MITRA, A., CHALLA, S., AND SIMICEVIC, N. Ultra-wideband signal propagation experiments in liquid media. *IEEE Transactions on Instrumentation and Measurement* 59, 1 (2010), 215–220.
- [36] TAO, B. Understand the mobile graphics processing unit. <http://www.embedded-computing.com/embedded-computing-design/understand-the-mobile-graphics-processing-unit>.
- [37] VON HIPPEL, A. R. *Dielectrics and waves*. Wiley, 1954.
- [38] WANG, G., QIAN, C., HAN, J., XI, W., DING, H., JIANG, Z., AND ZHAO, J. Verifiable smart packaging with passive rfid. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (2016), ACM, pp. 156–166.
- [39] WANG, J., XIONG, J., CHEN, X., JIANG, H., BALAN, R. K., AND FANG, D. Tagscan: Simultaneous target imaging and material identification with commodity rfid devices. In *Proc. ACM MobiCom* (2017), pp. 1–14.
- [40] WEI, T., AND ZHANG, X. Gyro in the air: tracking 3d orientation of batteryless internet-of-things. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking* (2016), ACM, pp. 55–68.
- [41] YUE, H., VIRGA, K. L., AND PRINCE, J. L. Dielectric constant and loss tangent measurement using a stripline fixture. *IEEE Transactions on Components, Packaging, and Manufacturing Technology: Part B* 21, 4 (1998), 441–446.
- [42] YUN, S., CHEN, Y.-C., ZHENG, H., QIU, L., AND MAO, W. Strata: Fine-grained acoustic-based device-free tracking. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services* (2017), ACM, pp. 15–28.