

The Lab Streaming Layer for Synchronized Multimodal Recording

Christian Kothe¹, Seyed Yahya Shirazi², Tristan Stenner³, David Medine⁴, Chadwick Boulay⁵, Matthew I. Grivich⁶, Tim Mullen¹, Arnaud Delorme², and Scott Makeig²

¹*Intheon Labs, San Diego, CA, United States*

²*Swartz Center for Computational Neuroscience, University of California San Diego, La Jolla, CA, United States*

³*Institute of Medical Psychology and Medical Sociology, University Medical Center Schleswig Holstein, Kiel University, Kiel, Germany*

⁴*Diademics Pty Ltd, Melbourne, Australia*

⁵*Ottawa Hospital Research Institute, Ottawa, Canada*

⁶*Neurobehavioral Systems, Berkeley, CA, United States*

Abstract—Accurately recording the interactions of humans or other organisms with their environment or other agents requires synchronized data access via multiple instruments, often running independently using different clocks. Active, hardware-mediated solutions are often infeasible or prohibitively costly to build and run across arbitrary collections of input systems. The Lab Streaming Layer (LSL) offers a software-based approach to synchronizing data streams based on per-sample time stamps and time synchronization across a common LAN. Built from the ground up for neurophysiological applications and designed for reliability, LSL offers zero-configuration functionality and accounts for network delays and jitters, making connection recovery, offset correction, and jitter compensation possible. These features ensure precise, continuous data recording, even in the face of interruptions. The LSL ecosystem has grown to support over 150 data acquisition device classes as of Feb 2024, and establishes interoperability with and among client software written in several programming languages, including C/C++, Python, MATLAB, Java, C#, JavaScript, Rust, and Julia. The resilience and versatility of LSL have made it a major data synchronization platform for multimodal human neurobehavioral recording and it is now supported by a wide range of software packages, including major stimulus presentation tools, real-time analysis packages, and brain-computer interfaces. Outside of basic science, research, and development, LSL has been used as a resilient and transparent backend in scenarios ranging from art installations to stage performances, interactive experiences, and commercial deployments. In neurobehavioral studies and other neuroscience applications, LSL facilitates the complex task of capturing organismal dynamics and environmental changes using multiple data streams at a common timebase while capturing time details for every data frame.

Index Terms—Brain/Behavior Quantification and Synchronization (BBQS), Multimodal recording, Mobile Brain/Body Recording (MoBI), Real-time synchronization.

I. INTRODUCTION

1
2 Recording and modeling brain dynamics supporting active, natural
3 cognition involving eye movements, motor and other behavior is
4 becoming an integral part of neurobiological research and requires
5 multimodal recording of the organism’s neural processes and interac-
6 tions along with concomitant changes in its environment. Successful
7 multimodal recording demands adequate temporal resolution and
8 precise synchronization of concurrently recorded data streams. In
9 human neuroscience, mobile brain/body imaging (MoBI) [1] is a
10 multimodal recording concept requiring synchronized recording of
11 brain, behavior, and environmental data streams with near millisecond
12 (msec) resolution. Maintaining synchronization at this scale between
13 brain (electro/magnetoencephalography, EEG/MEG; functional near-
14 infrared spectroscopy, fNIRS, etc.), behavioral (e.g., body motion
15 capture and eye tracking), physiological (electromyography, EMG,
16 etc.), and environmental data (video, treadmill, balance plate, robots,
17 or other agent positions and forces, sensory stimulation, etc.) often
18 requires multiple computer systems with no common, hardwired
19 clock to relate the timing of their outputs.

20 Here, we describe the Lab Streaming Layer (LSL), a software frame-
21 work that is helping researchers across academic and industrial settings
22 meet the challenge of multimodal recording through its ability to
23 collect data streaming from multiple devices and platforms operating
24 asynchronously on a local area network (LAN) along with msec-level
25 time synchronization and broad hardware and software compatibility.
26 LSL is a freely available open-source project under the umbrella of a

27 dedicated GitHub organization <https://github.com/labstreaminglayer>,
28 plus individual core repositories available from the Swartz Center
29 for Computational Neuroscience (SCCN) (meta-package and core
30 library). A listing of over 150 known LSL-compatible device classes
31 is compiled at <https://labstreaminglayer.org>, which also serves as a
32 landing page for finding tooling, documentation, and other resources.
33 LSL is supported by an active community of international contributors
34 (several of whom are among the coauthors), and at this point, two
35 annual workshops, one in Europe and one in the U.S., bring together
36 users, contributors, and developers, and present learning opportunities
37 for newcomers to the platform. Organizers currently include the
38 SCCN and teams at the University of Oldenburg and TU Berlin.
39 LSL’s popularity cannot be explained by any one of its features –
40 rather, a focus on ease of use and robustness, a distributed model
41 that allows for mixing and matching of multiple computers (desktop
42 or mobile) and software from multiple vendors and open-source
43 projects, likely contribute to its appeal, as does the broad platform
44 compatibility including most major programming languages and all
45 major desktop and mobile operating systems, along with built-in time
46 synchronization. Lastly, strong network effects owing to LSL’s large
47 ecosystem and installed base likely represent an additional factor for
48 its wide appeal.

49 One of LSL’s technical features is the synchronization of distributed
50 neuroscientific data streams based on a peer-to-peer protocol modeled
51 after the Network Time Protocol (NTP) as specified in RFC 5905[2]. A
52 closely related component is LSL’s decomposition of timing error into
53 three components: a constant, a slow-varying, and a noise component,
54 which are each addressed separately. Using these two approaches,

Lab Streaming Layer network

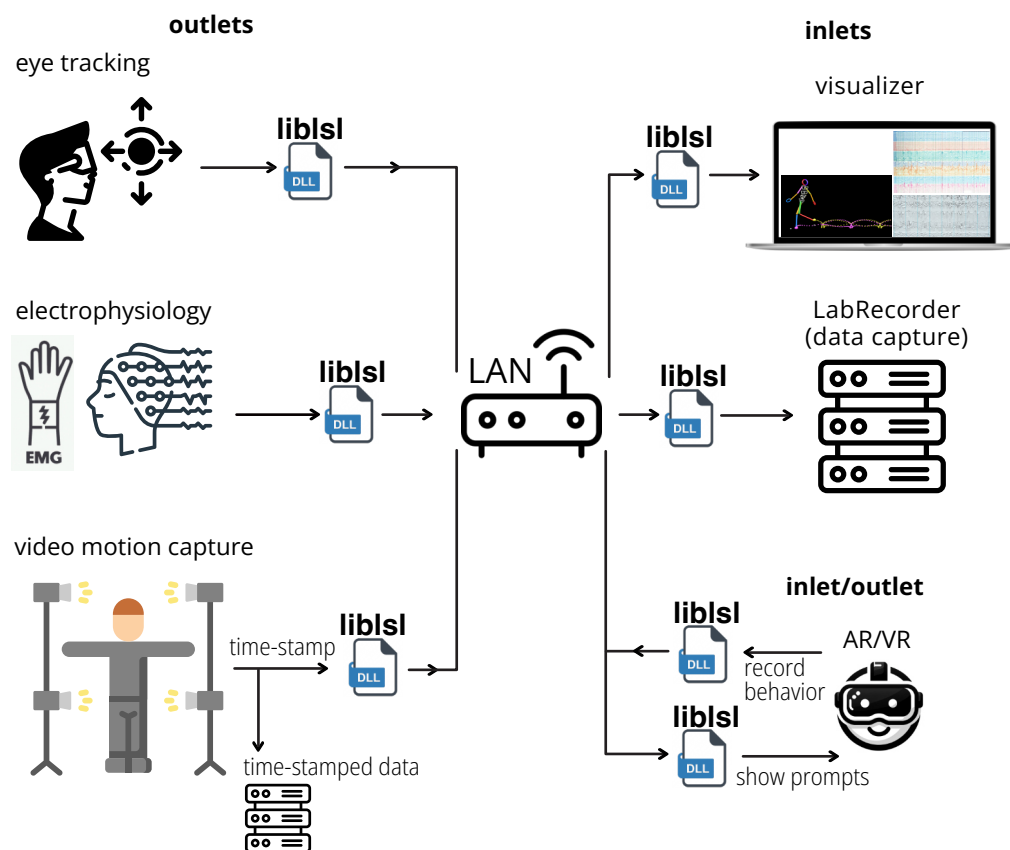


Fig. 1. **System overview.** The Lab Streaming Layer (LSL) creates a *network* connecting data acquisition, storage, and processing devices overlaying the local network (LAN) on which they are streamed. LSL handles publishing and subscribing to data streams, clock synchronization, accounting for network delays, and jitter using the LSL dynamic library (*liblsl*). LSL *outlets* publish data streams to the network that LSL *inlets* can subscribe to. *LabRecorder* is a space-efficient and high-throughput LSL recording program that can supervise recording of streams from any number of LSL *outlets*. Clients on the network include device integrations (seen on the left-hand-side), single- or multi-stream visualization or real-time processing components, and arbitrary stimulus presentation and response collection mechanisms.

1 LSL can ensure that timestamps associated with every data sample, 23
 2 collected across multiple acquisition devices and computers, are 24
 3 accurately compensated for intrinsic device delay, clock drift, and 25
 4 jitter, in the presence of variable network transmission latency. 26
 5 This capability is crucial in neuroscience research where near-msec 27
 6 precision can be essential for accurate data analysis and interpretation, 28
 7 particularly in studies involving complex brain/body dynamics, high- 29
 8 intensity biomechanics, and multi-subject interactions. 30

9 Challenges in collecting proper multimodal recordings include 1)
 10 the need to synchronize data streams from different platforms, 2)
 11 including data streams with heterogeneous sampling frequencies, 33
 12 3) set up and staff training of multiple recording workstations 34
 13 and (possibly proprietary) software, 4) interfacing with multiple 35
 14 proprietary data access APIs with limited OS and programming 36
 15 language support, documentation and learning resources, and 5) 37
 16 meeting challenges in data conversion, integration, storage, sharing, 38
 17 and reproducibility. Several hardware synchronization tools have been 39
 18 developed to address the pre-sampling synchronization in multimodal 40
 19 recordings. These include intricate systems of TTL (transistor- 41
 20 transistor logic) pulses, equipment for measuring throughput delays 42
 21 of recording instruments, and dedicating one instrument recording 43
 22 channel as a synchronizing clock [3]–[5]. 44

Recent advances in hardware-managed synchronization can improve common clock accuracy for digitally triggered events to tens of microseconds, including solutions based on shared clocks and analog-to-digital (A/D) converters and [6] radio-frequency trigger modules [7]. However, the use of hardware data synchronization approaches is very often not feasible in laboratories without resources to engineer special-purpose solutions across the range of proprietary acquisition systems researchers wish to use in their experiments. This is still more the case for low-cost and/or consumer-grade microelectronics-based systems that can now be used to record multimodal data inexpensively in paradigms, allowing, among others, greater degrees of participant mobility or at-home use.

Heterogeneous sampling frequency, platform inaccuracies, jitter, and sampling fluctuations make synchronization of the data stream using ‘start/stop’ events insufficient for neuroscience purposes. Such a setup may cause synchronization to drift by many milliseconds within mere minutes of data collection, which typically grows longer over longer recording durations. A recent study of multimodal MoBI data collection methods concluded that frequent TTL pulses are needed to retain millisecond synchronization between data streams [3]. Without this or some other hardware or software organizing method, data streams with different sampling frequencies typically

1 drift out of synchronization over time, compromising their worth for
2 joint analysis.

3 The setup and maintenance of professional timing equipment
4 across multiple workstations running mutually incompatible recording
5 software is time-consuming and may require a dedicated recording
6 technician and/or experimenter team to run, monitor, and document
7 the data collected by each system. A dedicated staff training process is
8 often required to learn to operate the acquisition software associated
9 with each system.

10 Finally, owing to the proprietary nature and variety of data collection
11 software and data access means for different systems and the need to
12 record metadata stored in different forms and locations, performing
13 data conversion and preprocessing, integration, annotation, storage,
14 analysis, and sharing is challenging. All these factors limit access
15 to high-quality research capabilities.

16 A. The Broader Landscape in Multimodal Recording

17 The LSL project was started in 2012 in response to an emergent
18 need for robust multi-modal data acquisition at Swartz Center for
19 Computational Neuroscience, UCSD by the first author (Christian
20 Kothe), where also the multimodal Mobile Brain/Body Imaging
21 (MobI) concept was originally proposed and first demonstrated [1].
22 Available software at the time for this purpose was a partly proprietary
23 package then in use at SCCN that followed a monolithic plugin-based
24 design and which was widely perceived as lacking reliability. Another
25 technology predating LSL is the Tobi Interface A [8], which aimed
26 to standardize the representation of biosignals, but which was also
27 implemented then in a monolithic manner. For robust distributed
28 simulator event tracking, an existing solution was HLA Evolved
29 [9], which influenced the attention paid to reliability. There was no
30 real-time data access protocol natively supported by multiple vendors
31 of EEG hardware, let alone a broader spectrum of neurobehavioral
32 modalities.

33 Since LSL is simultaneously a publish/subscribe overlay network
34 and API, a time-synchronization solution, a multi-modal time-series
35 and meta-data recording solution, and a real-time streaming tool
36 with native support for event data, there are to our knowledge
37 not many directly comparable alternatives. When reduced to its
38 network protocol aspect, some alternatives are ZeroMQ¹, MQTT²,
39 plain TCP/IP, and ³ (e.g., as used in BRAND [10]). In the audio
40 control domain an established protocol is Open Sound Control (OSC).
41 Besides Open Epyhs, another project supporting multiple types of
42 electrophysiology hardware is BrainFlow⁴, which currently supports
43 a range of low-cost and DIY devices. For instrument and lighting
44 control, respectively, well-known examples with good timing support
45 are MIDI and DMX, but these do not leverage existing Ethernet
46 or Wifi networking. However, it should be noted that even these
47 solutions can, and some have been, integrated with LSL via bridge
48 adapters. For time synchronization, alternatives are the precision time
49 protocol (PTP) [11], which however requires dedicated hardware,
50 and manual NTP-based synchronization. Without a doubt, numerous
51 research labs have developed countless pieces of in-house software
52 that acquires data from two or more devices, some of which are also

open source projects (e.g., Bonsai [12] with its focus on video and
electrophysiology analysis of behaving rodents mainly on Windows
workstations), but to our knowledge, none enjoy a degree of popularity,
broad plug-and-play device compatibility, and large installed-base as
LSL.

B. LSL Limitations

Despite the stringent LSL time synchronization guardrails described
below, LSL performance has some limitations. Most importantly, LSL
does not have access to any incoming data *until* the moment it is
received by the microprocessor (CPU) or microcontroller unit (MCU)
on which the LSL software communicating with the device is running.
Thus LSL cannot itself learn or estimate whatever *on-device delays*
within each recording device occurred (the intervals accruing between
data signal input and its arrival in the software). Measuring on-device
delay (or delay distribution) at least once for each acquisition stream
is therefore necessary to allow LSL to convert the recorded times of
data arrival into times of data capture. Once known, the delays, which
LSL models as constant in between setup changes, can be accounted
for and declared in software. This limitation is inherent to multimodal
neuroscience data acquisition systems engineered without common
hardware clock availability.

C. LSL Advantages

The LSL approach to synchronized aggregation of concurrent
data streams has three main advantages that together significantly
enhance the data acquisition process: 1) Facilitating multi-modal
data collection with heterogeneous and/or irregular sampling rates,
2) enabling distributed measurement and data processing across
multiple systems, and 3) streamlining both real-time and offline
access to time-stamped multimodal data through its companion *XDF*
file format.

The LSL unified Application Programming Interface (API) and pro-
tocol standardize data exchange across any number of measurement
modalities, creating a consistent real-time data stream access interface.
This simplifies initial device setup, allowing LSL-compatible clients to
require minimal or often no modifications to function with devices
from different vendors. The API also offers the flexibility to use
several of the most popular programming languages, allowing it to
be integrated into almost any piece of existing software with little
effort.

LSL allows time-synchronized stream readouts from all networked
devices, simplifying the experimental process to merely starting the
included recording devices and melding the received streams into an
integrated XDF data record using the LSL *LabRecorder* application (or
any equivalent of choice), eliminating the need to manage multiple data
file formats and increasing the efficiency of either near-real time or *post*
hoc data analysis. Moreover, LSL network protocol standardization
facilitates the distribution of data measurement and processing across
multiple computers without explicit network parameter configuration,
increasing data acquisition versatility.

II. SYSTEM OVERVIEW

LSL is a local network that runs on top of (or *overlays*) an Internet
Protocol (IP) network running at the experiment site. LSL network
peers can **publish** and **subscribe to** any number of **streams** of

¹<http://zeromq.org>

²<https://mqtt.org/>

³<https://redis.io>

⁴<https://brainflow.org>

1 single- or multi-channel time-series data (Figure 1). LSL regularly
 2 quantifies clock offsets (OFS) and round-trip time (RTT) between
 3 peers to enable data stream synchronization. Multi-channel samples
 4 of any stream published on LSL contain the channel values (of flexible
 5 type) and a time stamp assigned by LSL or the LSL integration ("LSL
 6 App") for the device. Peer access to LSL is set up using a dynamic
 7 library (*liblsl*) available for most POSIX-compatible platforms [13]
 8 including Windows, Linux, MacOS, Android, and iOS. The LSL API
 9 has been designed to "hide" the complexities of time synchronization
 10 and real-time network programming from both researchers and device
 11 manufacturers, while ensuring maximum network resiliency against
 12 dropped connections and data losses.

13 A. LSL Objectives

14 Chief goals governing LSL construction were: **a)** to simplify the
 15 discovery and selection of the published streams, **b)** to simplify
 16 publishing of active data streams to subscriber applications in
 17 near real-time, **c)** to supply sufficient metadata to allow for full
 18 interpretation of the transmitted time series, **d)** to solve the time-
 19 synchronization problem for concurrent data streams with an error low
 20 enough for most neurobehavioral research (i.e., at most msec-scale),
 21 **e)** to provide adequate out-of-the-box fault tolerance across a range
 22 of commonly-encountered failure scenarios (such as single-device
 23 failures, reconnects, restarts, intermittent network connectivity loss,
 24 and so forth), **f)** to establish a unified multimodal data representation,
 25 and **g)** to offer an API to access, transmit, and (when needed) store
 26 data from any set of data streams, regardless of modality.

27 Other possible objectives were explicitly *not* LSL design goals: *a)*
 28 building an online or *post hoc* data processing system (although such
 29 systems can easily be built on top of LSL), *b)* building an internet-
 30 scale and/or internet-facing data transport system, *c)* replacing or
 31 competing with existing data acquisition software (e.g., device drivers
 32 or applications), *d)* replacing or competing with non-signal intra-
 33 process or inter-process message queuing systems, or *e)* solving
 34 needs far outside physiological or neurobehavioral research (e.g.,
 35 high-energy physics).

36 B. LSL Design

37 The LSL software framework consists of three main components:
 38 the LSL API and language wrappers, the LSL core library (*liblsl*),
 39 and the LSL protocols (See Figure 2).

40 **The LSL API** is a unified interface to communicate with the LSL
 41 core library from external instruments and devices. To maximize
 42 compatibility and ensure a stable Application Binary Interface (ABI),
 43 LSL presents a C API in agreement with shared-library best practices,
 44 although the core is implemented in C++. Thanks to this stable ABI,
 45 support for other programming languages can be implemented with
 46 the C Foreign Function Interface (FFI), which enabled the creation
 47 of a wide range of wrappers for languages such as Java, C#, Python,
 48 Matlab, Rust, and several others. A header-only C++ API is also
 49 natively provided by the core library. These API wrappers provide
 50 the same metaphors, terminology, and functionality that the core
 51 C/C++ API provides.

52 Each existing API attempts to respect the idioms and standards
 53 of the language in which they are implemented. So, the Python API
 54 aims to be 'Pythonic' while the C API is an example of a 'classical'
 55 C style, while at the same time, all APIs cover an equivalent feature

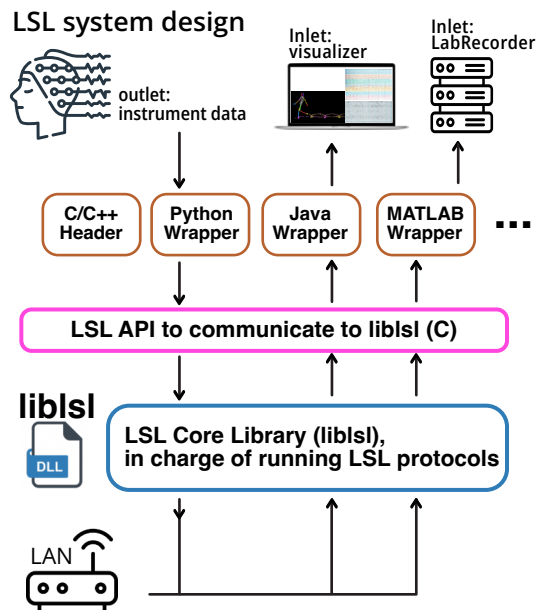


Fig. 2. **Lab Streaming Layer Design.** LSL consists of three main components: 1 LSL language wrappers and API, 2 LSL core library (*liblsl*), and 3 LSL protocols. The LSL API is a unified interface enabling communication with the LSL core library from external instruments and devices. The API was originally composed in C/C++ and is wrapped in other languages. The LSL core library (*liblsl*) is written in C++ and implements all features that LSL offers. The LSL protocols are the set of steps and standards required to establish reliable communication and synchronization between peers.

56 set. Developers can use the API to design executable programs to
 57 communicate with their peers on the network, publish data, and
 58 subscribe to streams from other peers.

59 A simple yet runnable example in Python that discovers, subscribes
 60 to, and then reads from an EEG stream on the LSL network is given
 61 in the following listing (equivalent examples are provided for all
 62 supported programming languages):

```
63
64 from pylsl import StreamInlet, resolve_stream
65
66 streams = resolve_stream('type', 'EEG')
67 inlet = StreamInlet(streams[0])
68
69 while True:
70     sample, timestamp = inlet.pull_sample()
71     print(timestamp, sample)
72
```

73 A corresponding simple example that generates 8 channels of
 74 random floating-point numbers and streams them to LSL at approx.
 75 200 Hz, here written in C++, is shown below. For best interoperability
 76 it is recommended to additionally specify meta-data such as channel
 77 labels, which is not shown here. Equivalent functionality is available
 78 for all other supported programming languages.

```
79
80 #include <chrono>
81 #include <lsl_cpp.h>
82 #include <thread>
83
84 const int nchannels = 8;
85
86 int main(int argc, char *argv[]) {
87     lsl::stream_info info("MyStream", "EEG",
```



```
1     nchannels, 200.0);
2     lsl::stream_outlet outlet(info);
3
4     float sample[nchannels];
5     while (1) {
6         for (int c = 0; c<nchannels; c++)
7             sample[c] = ((rand() % 1000) / 1000.0);
8         outlet.push_sample(sample);
9         std::this_thread::sleep_for(
10            std::chrono::milliseconds(5));
11     }
12     return 0;
13 }
```

15 **The LSL core library** (*liblsl*) is written in modern C++ and 73
16 manages features that LSL offers. Each peer needs to have a copy of 74
17 the *liblsl* to communicate with other peers on the network. Our effort 75
18 has been to maintain *liblsl* as a self-contained package to minimize 76
19 its dependencies on packages that are not shipped with the LSL 77
20 source code. Therefore, users should be able to compile the library 78
21 should the compiled code not be available on a given platform. 79

22 Internally, *liblsl* uses *pugixml* [14] for XML and XPath processing, 80
23 *loguru* [15] for logging with configurable verbosity and log targets, and 81
24 *Boost ASIO* [16], [17] for portable high-performance asynchronous 82
25 networking. 83

26 **LSL Network Protocols.** LSL internally implements five network 84
27 protocols to allow peers to create and maintain outlets to publish data 85
28 streams, inlets to subscribe to streams, and to stream information 86
29 objects each carrying all the requisite metadata for a data stream. By 87
30 protocols, we mean the steps and standards to establish outlets, inlets, 88
31 and metadata transfers. The five protocols are titled (1) Discovery, (2) 89
32 Subscription, (3) Stream transmission, (4) Metadata transmission, and 90
33 (5) Time synchronization. Adherence to the protocols is guaranteed 91
34 by the core library (*liblsl*). 92

35 1) *The Discovery Protocol:* The first stage in establishing 93
36 communication between inlets and outlets is stream discovery. An 94
37 application may discover outlet peers by broadcasting query messages 95
38 into the network via UDP broadcast and UDP multicast (RFC1112) 96
39 [18] to user-configurable multicast groups and awaiting responses. 97
40 The query message contains an XPath 1.0 [19] compliant query string 98
41 that specifies some metadata properties of the stream of interest (e.g., 99
42 type="EEG"). The host of each published stream on the network 100
43 will then respond to matching queries with a small response packet 101
44 that contains the essential properties necessary for establishing a 102
45 connection specific to the querying peer so that a single machine can 103
46 stream data to multiple peers at once. These include the name, type, 104
47 and unique identifier of the stream and are formatted as an XML 105
48 string. Responses to identical queries are cached for efficiency. 106

49 For convenience, all of this happens ‘under the hood’ of a single 107
50 LSL function call. The programmer of an LSL application need not 108
51 be concerned with the details of interfacing with a network stack 109
52 for all of this to work. Furthermore, queries can be transported over 110
53 several network protocols, including UDP broadcast and multicast 111
54 of various scopes, and can be done using IPv4 and/or IPv6. LSL 112
55 will correctly choose the right communication technique so that the 113
56 programmer can be agnostic of all the underlying network protocols. 114

57 The same LSL query protocol is used to automatically reconnect 115
58 to a peer should the connection be lost during a data transfer – for 116
59 example, if a software or network computer crashes, or a change 117
60 in network topology occurs. Connection recovery will be successful 118

61 even if the peer’s IP address has changed. This provides substantially
62 greater resilience than most protocols that cannot recover from a
63 change in IP addresses.

64 2) *The Subscription Protocol:* After a desired active outlet object
65 is discovered, the host application on the subscriber side will want
66 to connect a stream inlet to the outlet. This process is called an LSL
67 subscription, enacted by establishing a TCP connection to a network
68 endpoint advertised in response to the discovery query. A brief
69 two-way protocol negotiation handshake establishes this connection.
70 The handshake resembles HTTP/1.1 GET and its response [20].
71 The purpose of this handshake is to exchange several transmission
72 parameters such as the protocol version, byte order, buffer sizes,
support for floating-point subnormals, etc.

A mutually agreed-upon sequence of test-pattern data is also
transmitted to confirm that both parties can support the same protocol.
The metadata header (stream information object) is also transferred
from the host (outlet) to the client (inlet) to confirm that the endpoint
does carry the requested data stream. Once this exchange is completed,
the connection is formed, and time-series data will flow from the
outlet to the inlet until the connection is terminated.

3) *The Stream Transmission Protocol:* LSL transmits time-series
data as a byte stream split into packets by the underlying network layer.
Samples in the time series may be marked for immediate transmission
to enable use in real-time applications. This effectively indicates a
‘flush’ operation wherein the marked sample(s) are to be transmitted
as soon as the underlying network permits. The byte stream is a
sequence of encoded message frames. Every frame corresponds to
one sample and includes a losslessly delta-compressed timestamp
followed by the sequence of data values (bytes) encoded according to
the format agreed upon during the connection handshake. While the
underlying protocol is sample oriented, the choice between immediate
or deferred transmission allows users to send or receive time series
either sample-by-sample or at the granularity of multi-sample chunks,
where either side can choose to use either protocol, using easy-to-
use high-level functions (the above code listing shows sample-wise
sending and receiving).

4) *The Metadata Transmission Protocol:* In addition to time-
series data, a stream’s metadata must be transferred from peer to
peer. This metadata plays the same role as a file header in a time-
series recording and contains information such as the stream name,
type, channel count, sampling rate, etc. The metadata needs only be
transmitted once and is thus treated by LSL as ‘out-of-band’ data.
It is only transmitted on client request over a TCP connection. A
simple connection handshake also precedes this transfer.

The metadata is plaintext and structured in accordance with an
attribute-free subset of XML and can be of any length. The metadata
structure is not prescribed by LSL, but for interoperability it is
strongly recommended to adhere to a specification of content-types
(modalities such as EEG, Audio, Gaze, and so forth) and content
type specific nomenclature of XML fields. The latter specification
was co-developed with the XDF (extensible data format) project and
is available online from the XDF GitHub Wiki. Since this metadata
specification is plaintext XML, applications may extend and augment
this metadata in any way that is suitable for a given data stream
without breaking compatibility, or deviate when necessary.

5) *Time Synchronization Protocol:* A common use case of LSL
is streaming multimodal time series data from multiple peers to a
separate peer that subscribes (monitors and/or records) the multimodal

LSL local and network test for instrument delay

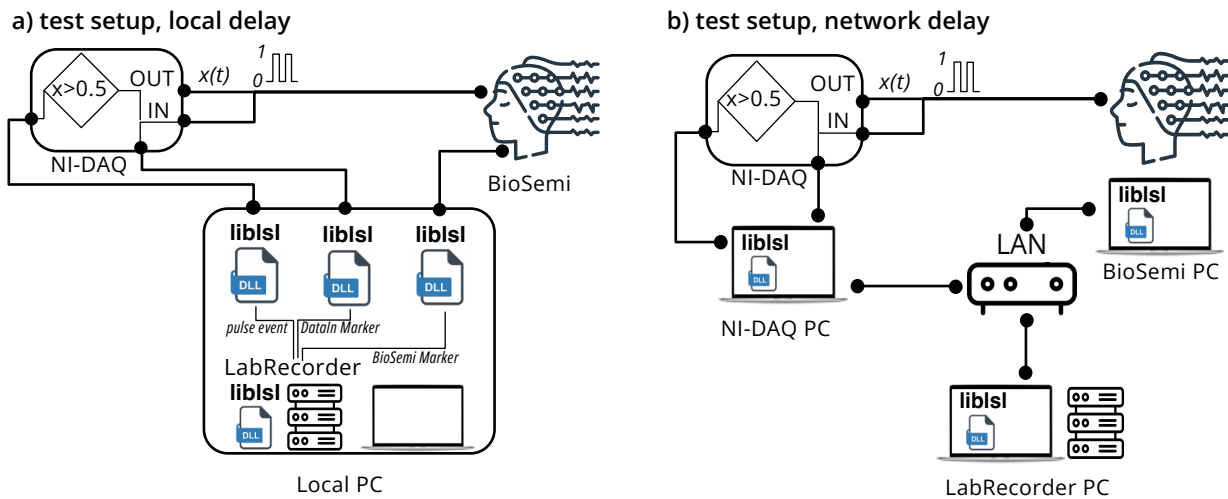


Fig. 3. **Synchronization performance setup.** The setup consists of a National Instruments Data Acquisition Box (NI-Daq) that generates a periodic pulse signal (DataOut) and receives the same signal (DataIn). The same NI-Daq is used to create an LSL marker when the pulse is going high. At the same time, a BioSemi Active-II receives the same pulse signal as an LSL stream. The BioSemi stream and the marker stream are recorded using LabRecorder, the native LSL recording program. The LSL marker stream is used to calculate the synchronization accuracy of the BioSemi stream. The local setup is using a single computer to connect to the NI-Daq and BioSemi devices and record the streams using LSL LabRecorder. The network setup is using separate computers to connect to the NI-Daq, BioSemi, and the LSL LabRecorder.

1 data. LSL's timestamping function returns the time of the most 24
 2 steady (i.e., monotonically increasing) high-precision computer clock 25
 3 available that has a minimum resolution of 1 msec or better (typically 26
 4 the machine uptime). The time offset between multiple computers' 27
 5 clocks, as well as their relative drift, is continually measured and 28
 6 accounted for by LSL when synchronization information is utilized.
 7 When an inlet peer wishes to synchronize its clock with the respective
 8 outlet peer, a structured packet exchange is initiated following the basic
 9 NTP model. Since clocks need to be periodically re-synchronized
 10 due to the drift, this process will be repeated regularly (e.g., by
 11 default, every 5 seconds). LSL employs the clock filter algorithm
 12 of the Network Time Protocol (NTP) [2] to account for random
 13 spikes in network transmission delay. This process uses multiple 29
 14 packet exchanges to estimate the clock offset (OFS) and round-trip 30
 15 times (RTT) between peers in rapid succession (e.g., ten times across 31
 16 200ms), yielding a set of OFSs and RTTs from which the one with 32
 17 the lowest RTT is retained. 33

Each packet exchange attempt for clock synchronization consists
 of a packet sent from the initiating peer to the receiver. This carries 34
 the local timestamp of the initiating peer and is noted as t_0 . The
 receiver then responds with two more timestamps, the receiving time 35
 of the original packet t_1 , and the time of resend t_2 . Upon receipt
 of this packet by the initiating peer, a final timestamp t_3 is taken. 36
 Then, 37

$$RTT = (t_3 - t_0) - (t_2 - t_1) \quad (1) \quad 40$$

$$OFS = ((t_1 - t_0) + (t_2 - t_3))/2 \quad (2) \quad 41$$

18 Therefore, RTT is the duration of the entire round trip minus the 43
 19 time spent on the receiving peer, and OFS is the averaged clock 44
 20 offset between the peers with symmetric network transmission delays 45
 21 canceled out. This measurement is a minimum-noise realization 46
 22 (because we choose the OFS at the minimum RTT) of the unbiased 47
 23 clock offset between the two peers. There can be a transmission time 48

asymmetry between the forward and backward network path (e.g.,
 due to driver implementation details), but the residual error after
 clock filtering is upper-bounded by the lowest delay of a machine's
 network implementation and is therefore assumed to be well under
 1 ms with most network hardware.

Using this time-varying measurement, LSL then constructs a model
 of the observed time stamps t_{obs} as a function of the time t_{actual}
 when the measurement actually occurred, an optionally smoothed
 estimate of the clock offset \overline{OFS} , a device-specific constant offset τ ,
 and a zero-mean noise term ε :

$$t_{obs} = t_{actual} + \tau + \overline{OFS} + \varepsilon \quad (3)$$

Using this formula, it is possible to recover t_{actual} for regularly
 sampled time series either using a recursive least-squares estimator
 in real time or linear regression in post-hoc data analysis, both of which
 are supported by LSL for the former and by XDF implementations
 for the latter.

C. The Extensible Data Format (XDF)

The Extensible Data Format (XDF) is an open-source and general-
 purpose natively multi-modal container format for multichannel time
 series data with extensive associated metadata. XDF is tailored towards
 biosignal data such as ExG, GSR, and MEG, but it can also handle
 data with a high sampling rate (like audio) or data with a high
 number of channels (like fMRI or raw video). In general, every
 data stream collected by the LabRecorder, along with metadata
 and synchronization information is recorded into a single XDF file.
 Crucially, XDF follows the policy of recording all timing-related
 ground-truth "as it happened", which allows for post-hoc analysis
 and recovery of data in case of misbehaving devices or intermittent
 failures during a recording. A result of this choice is that, while XDF
 importers present a simple interface similar to that of many other
 file importers, XDF files represent an exact record of what occurred

1 during an experiment, which can at times be complex, including a 57
2 device disappearing and later (e.g., after an unplanned battery swap) 58
3 reappearing. 59

4 In case of a high-bandwidth time series that may not be transferable 60
5 over the network (such as uncompressed video), each frame of the 61
6 stream may be timestamped and stored in the local machine (outlet) 62
7 while the timestamp information and the metadata would be sent 63
8 over LSL to the inlet machine and would be added to the XDF files. 64
9 Another scenario in which this may be favorable is when video data 65
10 falls under stricter privacy and regulatory requirements as personally 66
11 identifiable information (PII) than most other information that can 67
12 be recorded into an XDF file.

13 The XDF metadata is stored as XML content in an efficient 68
14 binary chunk-oriented container file format, and the recognized 69
15 metadata parameters are available at the XDF GitHub repository. 70
16 XDF predefines an extensible set of content-types (e.g., EEG, Audio, 71
17 NIRS, and so forth) and associated metadata specifications, following 72
18 a lightweight open process by which this specification is extended. 73
19 This allows a single file to maintain comprehensive yet extensible 74
20 modality-specific metadata on par with most unimodal biosignal file 75
21 formats. XDF tools are available for download at the XDF GitHub 76
22 page. A derived ANSI standard (ANSI/CTA-2060-2017) specifying 77
23 a file format for a consumer-grade variant of XDF has since been 78
24 published [21].

25 *D. Failure Resilience*

26 Preventing data loss is a major objective during data collection, 81
27 especially in multimodal data acquisition where the probability of 82
28 hardware issues grows linearly with the number of devices involved 83
29 in a given data collection setup. LSL is equipped with a number 84
30 of mechanisms for preventing catastrophic crashes and loss of data 85
31 to ensure smooth operation, even in the event of computer crashes 86
32 and lost network connections. To prevent data loss, LSL *outlet* and 87
33 *inlet* objects can use variable-size buffers that have a configurable, 88
34 arbitrarily large capacity. So, in case an *inlet* temporarily could not 89
35 receive data from an *outlet*, the data can be buffered until the *inlet* 90
36 can handle the transfer. The upper limit of all of this is the computer 91
37 resources and network throughput. 92

38 In the event of an *outlet* dropping out, any *inlets* connected to 93
39 the *outlet* will attempt to reconnect. An event will trigger within the 94
40 *inlet* to periodically search for the *outlet* and attempt to reconnect 95
41 as soon as the *outlet* is discovered. Since the *outlet*'s information 96
42 object can be created with a unique ID, this discovery will happen 97
43 automatically even if the *outlet* is recreated on a different computer 98
44 in the network and with a different IP address. 99

45 If an *outlet* drops out while an *inlet* is recording data, the 100
46 timing information for the dropped stream can be updated after the 101
47 rediscovery of the *outlet*, so that the outlet timestamp is consistent 102
48 with the timestamp information prior to the dropout. This behavior 103
49 is agnostic to the crash type and could resume recording of the 104
50 discovered *outlet* even if the disconnection is a result of changing 105
51 network topology, a computer crash, or hardware failure like a dead 106
52 battery. 107

53 Since these recovery processes happen automatically, the LSL user 108
54 is shielded from having to cope with anything other than potentially 109
55 a gap in a recorded data stream in the event that a device was 110
56 intermittently not recording data. XDF tools typically come with 111

built-in support for detection and correct handling of such data gaps.
These collective built-in efforts to recover connections between peers
realize LSL's failure resilience.

E. Software Stack

LSL includes an ecosystem of applications to publish and
subscribe to data streams, APIs in various languages built around
the core dynamic library (*liblsl*), an extensible data recording
format, *XDF*, post-hoc analysis for loading LSL synchronization
performance, and tools for performing offline time-synchronization.
This ecosystem can be accessed via the landing page and GitHub
organization and meta-repository. LSL also offers rich and open-
source documentation maintained by its developer community,
available at <https://labstreaminglayer.readthedocs.io>.

However, it is far beyond the scope of this article to do justice to
the greater LSL software ecosystem, which includes over a hundred
compatible client applications, some open source and others vendor-
native. Many applications in this greater ecosystem are hosted under an
umbrella GitHub organization, while many others are vendor-provided
data acquisition software with built-in LSL support, and an unknown
number of further LSL clients can be found via internet searches.
While this article focuses on acquisition devices, it is important
to note that the LSL ecosystem also includes a robust collection
of compatible stimulus presentation software, including most major
programs used for this purpose, which are indispensable for scientific
experimentation. Furthermore, the ecosystem includes software for
real-time processing of collected data (for example for brain-
computer interface or neurofeedback applications), visualization,
troubleshooting, experiment management, and various other tasks.

F. Continued Development and Maintenance

Researchers and programmers from both academic and commercial
sectors all over the world have contributed to the LSL source code
and APIs. However, changes to the core library (usually bug fixes)
are made very infrequently and with ultimate caution. Backward
compatibility with existing applications is maintained at all costs.
The bug rate is very low (less than one discovered every 6 months)
and so far, all bugs that were discovered were non-critical. Some
bugs seen so far include a few memory leaks and typing errors in
printing metadata and error messages. We have not found any bug
affecting the proper operation of sending and receiving data (the
primary LSL objective) in the past several years. Bugs in the LSL
application ecosystem and APIs are more common, but given the
stability and reliability of the core library and the simplicity of its
interface, these bugs are relatively trivial to identify and cannot affect
(i.e., crash) other LSL *inlets* and *outlets* – one of the less obvious
benefits of a decentralized design.

To maintain stability, unit tests covering a wide array of both internal
and API functions are run on all computing platforms for every change
committed to the source code. In addition, the library is periodically
stress-tested with hundreds of streams, randomized disconnects,
shutdowns, reconnects, and randomized stream parameters. During
such extreme network stress tests, some consumer-grade network
equipment has been found to be less reliable (i.e., crashing) than the
LSL implementation itself. Our dedicated benchmarks ensure that
changes in operating systems and libraries do not impair the data
exchange and synchronization performance.

instrument latency in the local LSL setup




Fig. 4. **Single-machine (local) synchronization performance.** The local setup is using a single computer to connect to the NI-Daq and BioSemi devices and record the streams using LSL LabRecorder.

III. TESTING AND RESULTS

LSL has been extensively tested and validated by the biosignal research community in several studies [6], [22]–[28]. Below, we provide some data concerning its performance on a local network (i.e., all LSL *inlets* and *outlets* running on a single machine), and on the distributed network synchronization performance. We provide a simple yet effective recipe to determine, for a given data instrument, the total delay of the data path for a given instrument, which is a sum of the internal hardware delay (e.g., on-device buffers), wireless transmission latency and operating system, device driver, and driver access latency, which we term in the following the "setup offset" τ .

Using a scientific grade analog-to-digital/digital-to-analog I/O device (National Instruments Data Acquisition Box, NI-Daq, Austin, TX) we created a periodic pulse signal (Figure 3). We used the same NI-Daq to receive the same signal (DataIn), and create an DataIn marker when the pulse was going high. To create the DataIn marker, we chose the time the recorded signal reaches halfway to its maximum amplitude. We also recorded the pulse event directly from NI-DAQ using LSL.

At the same time, we used another scientific-grade signal recording device (BioSemi Active-II, BioSemi B.V., Amsterdam, the Netherlands) and read the same pulse signal as an LSL stream. We used a similar threshold for the BioSemi-recorded pulse signal (i.e., halfway to maximum amplitude, BioSemi Marker), so that we could add time markers when the pulse signal went high. We recorded the BioSemi stream and the LSL marker stream using *LabRecorder*, the native LSL recording program.

Finally, we compared the timestamps of the marker stream and the 'high' points of the BioSemi stream. The NI-Daq data input stream was sampled at 10 kHz, and the BioSemi data stream was sampled at 2048 Hz.

We expected to observe a constant offset (setup offset) between the two markers (i.e., DataIn Marker and BioSemi Marker) due to the setup and network topology, plus some jitter. We ran the NI-Daq controller, BioSemi, and LabRecorder on (1) a single machine (Intel Windows 7) to test the LSL's local performance and (2) used separate network-attached machines for each of the NI-Daq controller, BioSemi, and LabRecorder (Intel Windows 7 for NI-Daq and Intel Windows 10 for each BioSemi and LabRecorder) to test LSL's network performance. We analyzed the difference of 1500

high-points generated by NI-Daq and BioSemi systems to quantify jitter and setup offset.

Here, we purposefully avoided using state-of-the-art machines because we wanted to test LSL performance on a more typical data acquisition setup.

A. Instrument Latency in a Local LSL Setup

The results showed a five-microsecond lead time between the time a DataIn Marker was issued and the *pulse events* satisfied our defined threshold (Figure 4a). This is well below the 100-microsecond resolution of the NI-Daq reader, so we considered this lead time negligible. Comparing the BioSemi Marker and the DataIn Marker latencies indicated a 12.20 ms setup offset between the two markers (Figure 4b). The jitter of this offset (i.e., the standard deviation of the lag (see Figure 4c) was 156 microseconds, below the ~500-microsecond Biosemi time resolution. Thus, the two streams could be aligned by removing this (pre-measured) device setup offset, and time jitter should not affect this alignment.

B. Instrument Latency in a Networked LSL Setup

To assess the setup offset of the instrument (in this example the BioSemi amplifier) in a distributed network, we separated the program controlling the NI-Daq (sending the DataIn Marker and storing *pulse events*), the program sending the BioSemi stream, and LabRecorder to network-attached computers. The results showed an even smaller setup offset between the DataIn Marker and the BioSemi Marker than the results observed in the single-machine LSL performance test (here, networked offset: 6.26 ms, vs. local offset: 12.20 ms, (Figure 5a). The offset jitter (presented as the standard deviation of the offset, (Figure 5b) was 145 microseconds, similar to the results from the local network experiment.

This offset decrease might have arisen from the separation, here, of the BioSemi and NI-Daq machines and potentially by faster performance of the BioSemi application and the associated driver running on Windows 10. However, the total setup delay for a given instrument is frequently dominated by device transmission delays, including large on-device buffer sizes that are only periodically transmitted, wireless (e.g., Bluetooth) protocol transmission latencies, and may add up to several 10s of milliseconds. Such discrepancies underpin the importance of testing setup offset (including device

instrument latency in the network-attached LSL setup

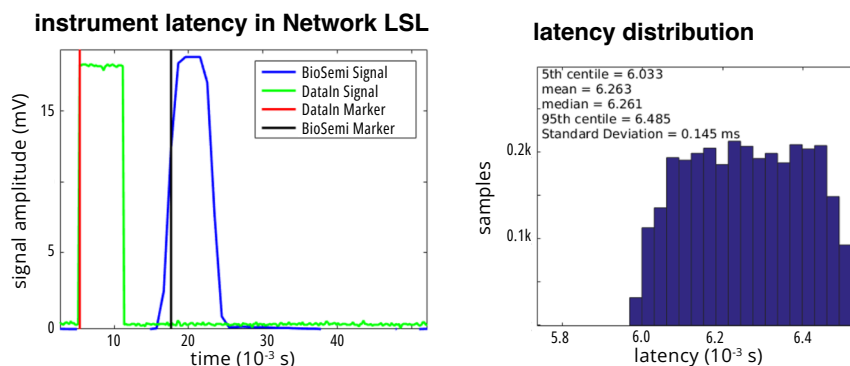


Fig. 5. **Network synchronization performance.** The network setup is using separate computers to connect to the NI-Daq, BioSemi, and the LSL LabRecorder.

1 throughput) for all devices and configurations before recording 41
2 experiment data. Setup offsets can be manually added to the 42
3 metadata while the other potential ad-hoc offsets caused by the 43
4 network delay or asynchrony would be recorded into the XDF 44
5 automatically. Both types of offsets will be addressed upon importing 45
6 the XDF files with the help of the LSL Time synchronization 46
7 protocol (II-B5) and using the `load_xdf.m` function (https://github.com/xdp-modules/xdp-Matlab/blob/master/load_xdf.m). 47
8

9 C. Determining the Setup Offset

10 As we demonstrated above, adjusting recording times for setup
11 offset is imperative for successful multimodal data acquisition and
12 synchronization. Modifying the setup configuration (e.g., moving an
13 *outlet* from one machine to another) may change the setup offset.
14 Any change in network configurations or updates to their software,
15 drivers, or operating systems should prompt a recheck. Here, we
16 present a simple yet effective algorithm to determine setup offset
17 for every instrument, a procedure similar to that described above in
18 III-A.

19 To determine the setup offset of an instrument, we suggest using a
20 microcontroller unit (e.g., an Arduino) board to send TTL pulses to
21 both the LSL network and to the instrument as a data input (Figure 6).
22 Publishing the TTL pulse as a DataIn Marker can be accomplished
23 through a control software that registers the TTL pulses, or can be
24 directly published by the MCU, since the LSL developer community
25 has provided support for running *liblsl* on some MCUs. The data
26 from the instrument should then be streamed to the LSL network.
27 Both the DataIn Marker and the instrument data should be recorded
28 using *LabRecorder*. The setup should be chosen in a way that most
29 exactly represents the experiment configuration. After reconstructing
30 a marker that corresponds to the TTL pulses from the instrument
31 data (instrument marker, similar to the BioSemi Marker in III-A),
32 the average offset between timestamps of the DataIn markers and
33 the instrument marker is the setup offset.

34 We should note that setup offset can be either positive or negative.
35 A positive offset means that the instrument marker occurs after
36 the DataIn marker, indicating an instrument lag in capturing and
37 transmitting the data to the recorder. A negative offset means the
38 instrument marker occurs *before* the DataIn marker; this may happen
39 for sensory triggers (e.g., auditory pulses) where the instrument
40 marker is the time that the trigger pulse is sent to the auditory 81

transducer (e.g., a loudspeaker), while the DataIn marker indicates
the time at which the transducer actually produces the pulse.

A successful setup with sub-millisecond internal delay using
an affordable MCU board (Arduino) has been benchmarked and
could be easily replicated from [29]. A commercial solution using
dedicated hardware for determining setup offsets is also available
from Neurobehavioral Systems, Inc.

48 Pitfalls and Tweaks

49 LSL can address some known hardware failures or network
50 connectivity issues. Sometimes, a hardware device may exhibit a
51 significant change in sampling rate (e.g., in our experience, a webcam
52 that frequently switches between 30 and 60 frames per second) or
53 suffer from high and variable packet loss (e.g., a Bluetooth device that
54 goes in and out of operational range). In these cases, the `load_xdf`'s
55 attempt to linearly smooth the timestamps will significantly (even
56 catastrophically) distort the data. This can be checked by comparing
57 the effective sampling rate as quantified by `load_xdf` (as the number
58 of samples divided by the recording length) with the sampling rate
59 reported in the device metadata. If these two sampling rates are
60 not close to each other, we suggest calling `load_xdf` with the
61 flag `'HandleJitterRemoval'` set to `false`. Oftentimes it is
62 possible to recover such recordings with some manual effort thanks
63 to XDF's policy to record all underlying ground-truth timing data.

64 A similar issue can arise by using LSL through a wireless local
65 area network (WLAN). If there are multiple streams on a heavily
66 utilized WLAN, the clock offset packet exchange can sometimes
67 overload the network and cause gaps in the data. In this case, it may
68 be appropriate to optimize the LSL configuration file for WLAN.
69 The recommended settings for WLANs are:

```
70 [tuning]  
71 TimeProbeMaxRTT = 0.100  
72 TimeProbeInterval = 0.010  
73 TimeProbeCount = 10  
74 TimeUpdateInterval = 0.25  
75 MulticastMinRTT = 1.0  
76 MulticastMaxRTT = 30
```

This text can be placed in a file called `lsl_api.cfg`. If this file
is in the same folder as the device's LSL application, these settings
would only be applied to the device. If the file is in `~/lsl_api/`,

suggested procedure for **Setup Offset** determination

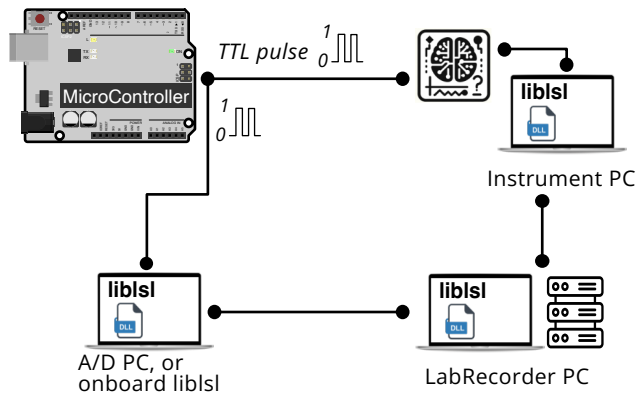


Fig. 6. **Setup offset determination algorithm.** The setup offset can be determined by sending a TTL pulse from a microcontroller board to the LSL network and to the instrument. The instrument data would be streamed to LSL, and the LSL marker would be recorded using LabRecorder. The setup offset would be the average offset between the DataIn marker and the instrument marker.

computational load of processing multiple streams across separate network-attached machines can at times outperform the setup offset (and latency) achieved by capturing all data streams on a single, perhaps heavily loaded, machine, which is made trivial thanks to LSL's ability to seamlessly discover streams across the network without additional configuration.

LSL as a purely software-based approach has an inherent limitation when no hardware triggering mechanisms are used, which is that LSL as a network is not aware of any latency occurring *within* the acquisition device or in the device drivers before data reaches the LSL application for the device. While LSL integrations can make reasonable assumptions, and some do, any residual offset in this latency, which typically amounts to a few 10s of milliseconds should be ascertained prior to conducting a study, ideally through testing using the actual devices and parameter settings to be used during subsequent recordings. A similar limitation applies to event marker time stamps pertaining to button presses or on-screen presentation, where again it is recommended to measure the input and/or display latency using off-the-shelf tools such as photodiodes or high frame rate cameras. Lastly, when the consistency of device sampling rate itself and/or the stability of its setup offset cannot be trusted, it may be necessary to implement a hardware-based data timing device to monitor the process. Therefore, while LSL can recover lost connections and compensate for offsets and jitter, an appropriate initial setup of the instruments and measuring setup offset are imperative for an optimally synchronized multimodal recording.

While LSL accommodates a relatively large buffer to minimize data loss in case of a connection drop or subpar network speed, given a long enough (e.g., a few minutes) network disconnection, the buffer may eventually run out with the resulting loss of data. Similarly, LSL data throughput is limited by network and computer capacity. While many data streams can be easily transferred at multiple KHz rates, some data streams, such as high-definition video, may saturate the bandwidth. In such a case, using lightweight compression before broadcasting the stream or storing the timestamped data on the local machine and only streaming the timestamps through LSL may resolve this issue.

A large ecosystem, transparent codebase and development, zero-configuration, excellent latency management, and reliability have made LSL a go-to solution for synchronized multimodal quantification of brain and behavior. Researchers can enjoy LSL with minimal and one-time initial setup and be sure that LSL will stream and store their multimodal data streams accurately and reliably. Finally, LSL development thrives on an open and welcoming community of enthusiasts. Anyone can join this effort via LSL's community hubs.

ACKNOWLEDGMENTS

The first version of the LSL software was written at the Swartz Center for Computational Neuroscience, UCSD, funded by the Army Research Laboratory under Cooperative Agreement Number W911NF-10-2-0022 as well as NINDS grant R01NS047293, and by a gift to UCSD from The Swartz Foundation (Old Field, NY).

CK and TM have received compensation from Intheon, which offers products and services that make use of LSL. TS, DM, CB, and MG have provided consulting services or have worked on products that use LSL.

REFERENCES

- [1] S. Makeig, K. Gramann, T.-P. Jung, T. J. Sejnowski, and H. Poizner, "Linking brain, mind and behavior," *Int. J. Psychophysiol.*, vol. 73, pp. 95–100, Aug. 2009.

the changes would be applied to the user globally. If the file is placed in an /etc folder (C:\etc on Windows), the tweaks will be global for all users.

Since applications can supply their own time stamps upon submitting a sample to LSL, potentially outside of the control of the user, it is possible to selectively ignore such time stamps via the user-facing configuration file. This can be necessary when a third-party application uses non-standard time stamps (e.g., from an alternative clock source such as on-device clocks). Since LSL tracks time offset between host machines and not between arbitrary application-chosen clocks, in such cases the recorded data would appear mutually unsynchronized. To enable this feature, the user can put the following lines into their `lsl_api.cfg`:

```
[tuning]
ForceDefaultTimestamps = 1
```

IV. SUMMARY AND CONCLUSION

The Lab Streaming Layer is a now well-established, reliable and easy-to-use multimodal signal acquisition, transmission, and recording platform tuned for synchronously recording multimodal brain and behavioral data. Oftentimes, using LSL with a given device can be as simple as enabling LSL support in a vendor-provided data acquisition software, if supported, or alternatively using one of the existing open-source integrations for the device, and recording the data on the same or another machine with the *LabRecorder* or another LSL-compatible recording tool. However, LSL also scales to complex setups involving multiple machines and several dozen acquisition devices or data streams. In one multiperson, multiple touchscreen simulation [30], we successfully used LSL to record from over 40 LSL data streams⁵ in recording sessions lasting multiple hours.

Our exemplar tests support the excellent sub-millisecond accuracy of the LSL timestamps. As our tests also showed, distributing the

⁵Two concurrent subjects, each with instruments including a 267-channel BioSemi, microphone, force plate, eye-tracking, three cameras, motion capture, and event marker streams.

- 1 [2] J. Martin, J. Burbank, W. Kasch, and D. L. Mills, "RFC 5905: Network time 74
2 protocol version 4: Protocol and algorithms specification." <https://datatracker.ietf.org/doc/rfc5905/>, June 2010. Accessed: 2023-11-27. 76
- 3 [3] F. Artoni, A. Barsotti, E. Guanziroli, S. Micera, A. Landi, and F. Molteni, "Effective 77
4 synchronization of EEG and EMG for mobile brain/body imaging in clinical 78
5 settings," *Front. Hum. Neurosci.*, vol. 11, p. 652, 2017.
- 6 [4] C. Maidhof, T. Kästner, and T. Makkonen, "Combining EEG, MIDI, and motion
7 capture techniques for investigating musical performance," *Behav. Res. Methods*,
8 vol. 46, pp. 185–195, Mar. 2014.
- 9 [5] D. Bannach, O. Amft, and P. Lukowicz, "Automatic event-based synchronization
10 of multimodal data streams from wearable and ambient sensors," in *Lecture Notes
11 in Computer Science*, vol. 135 of *Lecture notes in computer science*, pp. 135–148,
12 Berlin, Heidelberg: Springer Berlin Heidelberg, 2009.
- 13 [6] C.-H. Chuang, S.-W. Lu, Y.-P. Chao, P.-H. Peng, H.-C. Hsu, C.-C. Hung, C.-L.
14 Chang, and T.-P. Jung, "Near-zero phase-lag hyperscanning in a novel wireless
15 EEG system," *J. Neural Eng.*, vol. 18, Nov. 2021.
- 16 [7] G. L. Cerone, A. Giangrande, M. Ghislieri, M. Gazzoni, H. Piitulainen, and
17 A. Botter, "Design and validation of a wireless body sensor network for integrated
18 EEG and HD-sEMG acquisitions," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 30,
19 pp. 61–71, Jan. 2022.
- 20 [8] C. Breitwieser and C. Eibel, "TiA – documentation of TOBI interface a," *arXiv
21 [cs.NI]*, Mar. 2011.
- 22 [9] B. Möller, K. L. Morse, and M. Lightner, "HLA evolved – a summary of major
23 technical improvements," 2008.
- 24 [10] Y. H. Ali, K. Bodkin, M. Rigotti-Thompson, K. Patel, N. S. Card, B. Bhaduri, S. R.
25 Nason-Tomaszewski, D. M. Mifsud, X. Hou, C. Nicolas, S. Allcroft, L. R. Hochberg,
26 N. A. Yong, S. D. Stavisky, L. E. Miller, D. M. Brandman, and C. Pandarinath,
27 "BRAND: A platform for closed-loop experiments with deep network models,"
28 *bioRxiv*, p. 2023.08.08.552473, Aug. 2023.
- 29 [11] IEEE SA Standards Board, "IEEE standard for a precision clock synchronization
30 protocol for networked measurement and control systems," June 2020.
- 31 [12] G. Lopes, N. Bonacchi, J. Frazão, J. P. Neto, B. V. Atallah, S. Soares, L. Moreira,
32 S. Matias, P. M. Itskov, P. A. Correia, R. E. Medina, L. Calcaterra, E. Dreosti,
33 J. J. Paton, and A. R. Kampff, "Bonsai: an event-based framework for processing
34 and controlling data streams," *Front. Neuroinform.*, vol. 9, Apr. 2015.
- 35 [13] "IEEE standard for information technology–portable operating system interface
36 (POSIX(TM)) base specifications, issue 7," 2018.
- 37 [14] A. Kapoulkine, "pugixml: Light-weight, simple and fast XML parser for c++ with
38 XPath support."
- 39 [15] "loguru: Python logging made (stupidly) simple."
- 40 [16] C. Kohlhoff, "Boost.asio - 1.82.0." https://www.boost.org/doc/libs/1_82_0/doc/html/boost_asio.html. Accessed: 2023-7-1.
- 41 [17] S. Koranne, "Boost c++ libraries," in *Handbook of Open Source Tools*, pp. 127–143,
42 Boston, MA: Springer US, 2011.
- 43 [18] D. S. E. Deering, "Host extensions for IP multicasting." RFC 1112, Aug. 1989.
- 44 [19] "XML path language (XPath)." <https://www.w3.org/TR/1999/REC-xpath-19991116/>. Accessed: 2023-7-13.
- 45 [20] R. T. Fielding, M. Nottingham, and J. Reschke, "RFC 9110: HTTP semantics."
46 <https://www.rfc-editor.org/rfc/rfc9110.html>. Accessed: 2023-7-13.
- 47 [21] American National Standards Institute, "Standard for consumer EEG file format,"
48 Nov. 2017.
- 49 [22] S. Bustamante, J. Peters, B. Scholkopf, M. Grosse-Wentrup, and V. Jayaram,
50 "ArmSym: A virtual human–robot interaction laboratory for assistive robotics,"
51 *IEEE Trans. Hum. Mach. Syst.*, vol. 51, pp. 568–577, Dec. 2021.
- 52 [23] T. Kang and C. Wallraven, "Gotta go fast: Measuring input/output latencies of
53 virtual reality 3D engines for cognitive experiments," *arXiv [cs.HC]*, June 2023.
- 54 [24] D. Weber, S. Hertweck, H. Alwanni, L. D. J. Fiederer, X. Wang, F. Unruh,
55 M. Fischbach, M. E. Latoschik, and T. Ball, "A structured approach to test the
56 signal quality of electroencephalography measurements during use of head-mounted
57 displays for virtual reality applications," *Front. Neurosci.*, vol. 15, p. 733673, Nov.
58 2021.
- 59 [25] J. Levitt, Z. Yang, S. D. Williams, S. E. Lütsch Espinosa, A. Garcia-Casal, and L. D.
60 Lewis, "EEG-LLAMAS: an open source, low latency, EEG-fMRI neurofeedback
61 platform," *bioRxiv*, Nov. 2022.
- 62 [26] S. Iwama, M. Takemi, R. Eguchi, R. Hirose, M. Morishige, and others, "Two
63 common issues in synchronized multimodal recordings with EEG: Jitter and
64 latency," *bioRxiv*, 2022.
- 65 [27] M. Merino-Monge, A. J. Molina-Cantero, J. A. Castro-Garcia, and I. M. Gomez-
66 Gonzalez, "An easy-to-use multi-source recording and synchronization software
67 for experimental trials," *IEEE Access*, vol. 8, pp. 200618–200634, 2020.
- 68 [28] S. Blum, D. Hölle, M. G. Bleichner, and S. Debener, "Pocketable labs for everyone:
69 Synchronized multi-sensor data streaming and recording on smartphones with the
70 lab streaming layer," *Sensors (Basel)*, vol. 21, p. 8135, Dec. 2021.
- 71 [29] S. Appelhoff and T. Stenner, "In COM we trust: Feasibility of USB-based event
72 marking," *Behav. Res. Methods*, vol. 53, pp. 2450–2455, Dec. 2021.
- 73 [30] C. Kothe, T. Mullen, and S. Makeig, "Strum: A new dataset for neuroergonomics
research," in *2018 IEEE International Conference on Systems, Man, and
Cybernetics*, pp. 77–82, IEEE, Oct. 2018.