

Audiosockets: A Python Socket Package for Real-Time Audio Processing

Nicolas Shu* and David V. Anderson

Department of Electrical and Computer Engineering,
Georgia Institute of Technology

***Corresponding Author**

Nicolas Shu, Department of Electrical and Computer Engineering, Georgia Institute of Technology.

Submitted: 2024, May 13; Accepted: 2024, Jun 12; Published: 2024, Jul 26

Citation: Shu, N., Anderson, D. V. (2024). Audiosockets: A Python Socket Package for Real-Time Audio Processing. *J Electrical Electron Eng*, 3(4), 01-04.

Abstract

There are many packages in Python which allow one to perform real-time processing on audio data. Unfortunately, due to the synchronous nature of the language, there lacks a framework which allows for distributed parallel processing of the data without requiring a large programming overhead and in which the data acquisition is not blocked by subsequent processing operations. This work improves on packages used for audio data collection with a light-weight backend and a simple interface that allows for distributed processing through a socket-based structure. This is intended for real-time audio machine learning and data processing in Python with a quick deployment of multiple parallel operations on the same data, allowing users to spend less time debugging and more time developing. *Index Terms:* python, audio, speech, socket programming, real-time processing

1. Introduction

Machine learning and analysis of audio signals is a rapidly growing field. Although most of the machine learning research for audio processing is done offline, practical systems often require on-line, real-time operation. There are many packages available in Python for audio loading and a few which allow one to record audio in real-time [1-5]. However, often it is the case in a synchronous language like Python, where a process that is continuously recording audio may only record, and cannot asynchronously spend its compute time processing the audio through an algorithm (e.g. a neural network). There exist solutions such as using Robot Operating System (ROS) to perform message passing between parallel processes; but, although easy to use and very useful for multiple tasks, ROS is a relatively large software which contains multiple features that may be too heavyweight for a simpler task which only requires message passing [6]. Currently, there lacks an off-the-shelf software package for Python which is simple to use and quickly deployable for real-time processing for Python-related audio projects. This work introduces audio sockets, a socket-based package which uses the Sound device package to record audio, deploys a local server and allows for multiple clients to behave as processors for the incoming audio [5]. This package handles all of the socket programming in the backend and is designed to minimize the user's efforts on designing the message passing, and maximize their time developing new algorithms.

1.1 Related Work

While there are packages specifically centered on to do audio processing such as PySox and Librosa, there are a few options

which one may take to perform real-time inference over audio data. PyAudio is an example of a package which may be used for real-time inference. It allows a user to open audio streams and contains a callback mode, which allows a user to place the data, which is stored as byte strings, in a queue without blocking the script, making it a very powerful tool for a programmer. The one minor downside of PyAudio is that it requires a good understanding of how microphones operate, and it requires a significant amount of programming and attention to detail before getting the code to function properly. Furthermore, the PyAudio package by itself does not allow one to distribute the post-processing work across different threads.

PyJack is another package which interfaces with the JACK (JACK Audio Connection Kit) API, however, the code was originally written for Python 2.7, and, at the time of this writing, the latest release was in March 2011, leaving many of the connections deprecated. PyGame is a solution that has a similar interface as PyAudio, where it requires one to create an object which will use a callback function to append chunks of data in bytes format to a list, but it also suffers from the lack of work distribution across different threads for post-processing the data. Sound device is another feasible solution as it is also capable of using callback functions, allowing one to store the data in a queue, but it also already converts the output data as numerical values, ready for other use cases.

2. Method

This package uses socket programming as its backend communication protocol and the package Sound device to obtain

streams of audio from a microphone as well as NumPy [7]. Therefore, the only packages needed to be installed are sound device and NumPy, as well as Port Audio for your operating system. The architecture for this package is shown in Figure 1.

This package deploys a local server on one's system, which will bind a socket to an IP address and a pre-defined port, and will start to listen for new connections from clients. Given that a local area network exists, this package is capable of sending data over a network. Clients will have two different types: "recorder" and "processor", and the server will know how to handle each of those two types of clients. The two types of clients are seen as equal to the server, and through Python's threading module, there is a thread which waits for clients to connect to it while another thread starts to connect a new incoming connection.

For the case of a recorder, once initialized, the recording node will create a socket and will attempt to connect to the predefined address of the server. The node will attempt to connect to the server at 1Hz, and once it does connect, it will create a queue object and create a stream object which continuously dumps the audio data from the microphone in the queue. Periodically, it will start to send the audio data to the server. To do so, it will perform two major steps. The first step is to serialize the data and to determine the number of bytes required to send the message to the server. Once the number of bytes has been determined, the node will create a header message of size of 64 bytes, and will send a message header to the server specifying the number of bytes for the full audio data that the server is expected to receive.

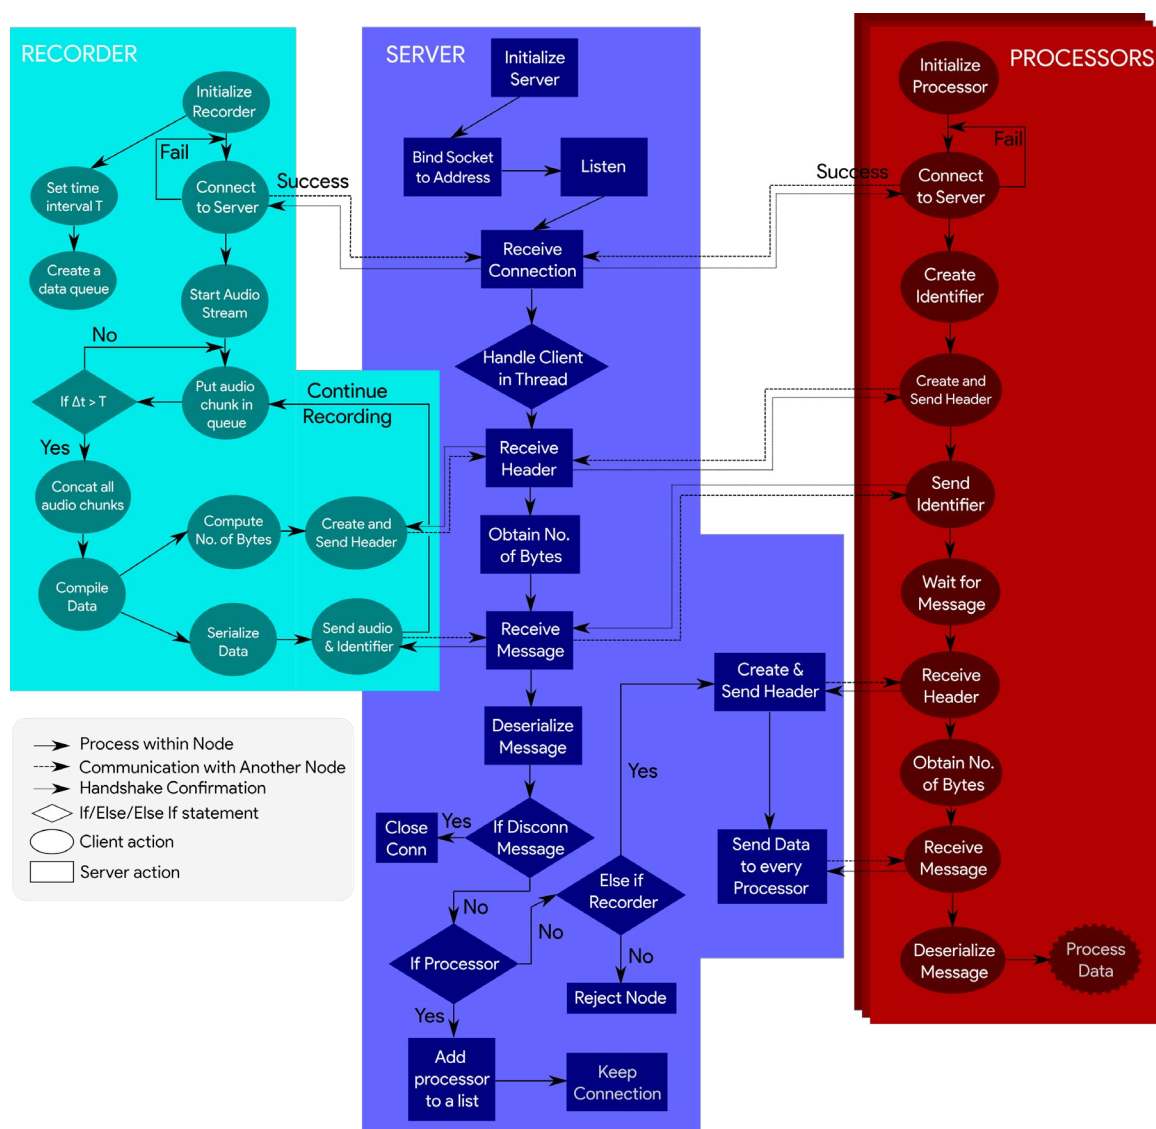


Figure 1: Mechanism for the Audiosockets Package. The Different Colored Superblocks Represent the Different Parallel Processes Which are Operating and Communicating with Each Other. The Superblocks Containing Circular Processes (i.e. Recorder and Processors) are the clients, and the Superblock with Rectangular Processes Represents the Local Server Which Is to Be Deployed in a System.

Once the server receives a new connection from the Recorder client, and it will wait for a new messages to be received. It will expect two messages. First, a message of size 64 bytes which specifies the number of bytes of the actual message to be received, while the second will be the actual message. Once the server receives the header, it will send back a message to the client stating that it has successfully received the header message, and it's ready to receive the data. The recorder will then send a message containing the all the data including: the audio data, the sampling frequency, the current time that the message was sent, and a type identifier for the server so that the server knows how to handle the given client. The server will then receive the data message, and send back a confirmation back to the client stating that it has received the data message. Since the stream of audio from Sound device works in parallel to the script, it will continue to gather more data, and then the recorder will repeat the recording steps until it is time to send a new message to the server.

Unfortunately, in socket programming, just because a connection expects to receive a specific sized message, let us say 2048 bytes, it does not mean that it will receive the full message at once. Therefore, as the server receives the stream of data, it will continue to check for the rest of the fragments of the message until it receives the entire message, and only then will it send the confirmation back to the client that it has received the data message. At this point, the server first needs to deserialize the message onto an actual object. Once the object has been interpreted, the server will check the type of client that just sent the message. If the client was a recorder, it will follow a recording protocol. If the client was a processor, then it will follow a processor protocol. But if it was none of the above, it will send an error message back to the client and dump the data.

If the server receives a data message from a processor, the first

thing it does is to register the processor to the server by name on a pointer, and it holds that connection opened. If the server receives a data message from a recorder, then it sends the data to all of the registered processor connections.

The processor client behaves in a very simple manner. Once it is initialized, as described before as a characteristic of it being a client, it attempts to connect to the server. Once a connection is established, the processor node sends a header message specifying the size of the data message that it will send the server, and once the server has confirmed that it has received the message, the processor node will send a data message containing a name identifier and a type identifier. As described before, the processor node will be registered to the server as a processor, and then it will block itself waiting for a header message from the server. Once it receives a header message from the server, it will send a confirmation back to the server that it has received the header message, and then will wait to receive a data message, which will, in turn, be deserialized and interpreted by the object. The object will then have a method that will process the data in whichever way the user desires.

To deal with disconnections, all of the clients have methods that will send disconnecting data messages to the server. If the server receives a disconnecting message from a client, it will delete the processor from the registered processors (if the client is a processor), the server will send a confirmation message back to the client that it will disconnect the client, and then proceed to close the connection between the server and the client.

The versatility of this method allows for one to connect to multiple processors at once and to disconnect any of the clients (processors or recorder) at any point in time, and reconnect it again, as shown in Figure 2

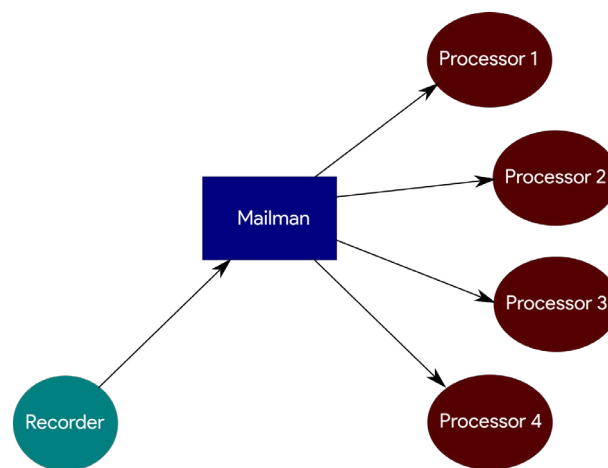


Figure 2: A Simplistic Overview Diagram of the Capabilities of the Package.

3. Usage

Here, a baseline is described on how to use this package.

3.1 Network Setup

In order to setup the network, one needs to create a JSON file which contains a few parameters

• **“SERVER”**: The string IP address of the machine where server

is deployed to

- **“PORT”**: The port number in the local computer where the server will be deployed and the port where the clients will try to bind to.
- **“HEADER”**: The number of bytes defined for size of the header message
- **“FORMAT”**: The type of which text data will be encoded or

decoded

- **“DISCONNECT_MSG”**: The disconnecting message which the clients will send to the server and the message that the server will be expecting to disconnect the client
- **“logging_format”**: The format for the logging module to report the logs
- **“logging_level”**: The level of logging to be displayed

This an example is as shown below

```
{
  "SERVER": "172.16.12.10",
  "PORT": 5050,
  "HEADER": 64,
  "FORMAT": "utf-8",
  "DISCONNECT_MSG": "DISCONNECT",
  "logging_format":
    "%(asctime)s-%(message)s",
  "logging_level": "info"
}
```

3.2 Server Deployment

In order to deploy a server, one needs to import the MailmanSocket, pass the path to the JSON file, and run the start() method. An example is shown below:

```
from audiosockets import MailmanSocket
if __name__ == "__main__":
    mailman = MailmanSocket(
        "server_info.json")
    mailman.start()
```

3.3 Recorder Deployment

In order to deploy the recorder, one needs to import the RecorderSocket, pass the path to the JSON file specified above, and run the start() method. An example is shown below

```
from audiosockets import RecorderSocket
if __name__ == "__main__":
    recorder = RecorderSocket(
        "server_info.json")
    recorder.start()
```

3.4 Processors Deployment

The processor is slightly more involved, but that is up to the user with respect to how complex does he or she wishes to process the data. First, one needs to import the class BaseProcessorSocket, and inherit from it. One must pass the arguments to the parent's `__init__()` method, and then one must overwrite the method `process_data(data)` method, which takes in a single argument, with anything the user wishes to do with the data. The example below shows the usage of a LogMelSpectrogram class that computes the log Mel spectrogram of a signal given a sampling frequency. Finally, one must run the method `start()` so that it binds to the server and becomes part of the network. [4]

```
from audiosockets import BaseProcessorSocket
from audiosockets.utils import \
    LogMelSpectrogram as LMSProcessor
class LMSProcessor(BaseProcessorSocket):
    def __init__(self,*args, **kwargs):
        super().__init__(*args, **kwargs)
    def process_data(self,data):
        fs = data["fs"]
        audio = data["data"]
        lms = LogMelSpectrogram(fs) (audio)
        print(lms.shape)
if __name__ == "__main__":
```

```
processor = LMSProcessor(
    "VAD",
    "server_info.json")
processor.start()
```

4. Conclusion

Before this work, there lacked an off-the-shelf open-source framework that would allow one to quickly deploy real-time audio processing in Python. For some solutions, a lot of understanding with regards to lower-level analog-to-digital processing was required. Other solutions would hit a roadblock due to the synchronous behavior of Python, where, for a real-time application, if an algorithm takes longer to process the data than the audio collection itself, it then becomes the bottleneck to the pipeline, and data may potentially be lost. Furthermore, they would all require additional work to get a distributed computing framework where different nodes would contain different algorithms, and would not bottleneck the node collecting the data. This work fills all of the gaps previously mentioned. It allows one to quickly deploy a real-time audio processing solution, where a node is responsible for only collecting data and sending it to a local server, and the local server's role is to simply receive any data the first node first sent it and sends it to other nodes which each may contain an algorithm which likely does not require to be re-initialized. As an example, neural networks require, relatively speaking, a long time to be initialized and prepared for inference. Depending on the capabilities of the machine, this framework allows the user to deploy several different neural networks on their GPU, each of which may be specialized for a different task, and the local server sends the data to each of them. At no point is there a loss of data, and with a minimal amount of work, a fully functional real-time processor can be deployed.

Acknowledgements

We like to thank Mouhyemen Khan for the fruitful discussions over ideas which could be used in order to implement and achieve the goal of this work.

References

1. “Pham, H. (2023). PyAudio: Cross-platform audio I/O for Python, with PortAudio.”
2. Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., & Van Mulbregt, P. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature methods*, 17(3), 261-272.
3. McFee, B., Raffel, C., Liang, D., Ellis, D. P., McVicar, M., Battenberg, E., & Nieto, O. (2015). librosa: Audio and music signal analysis in python. In *SciPy* (pp. 18-24).
4. Bastian Bechtold, “Pysoundfile.”
5. M. Geier. (2020). “Sounddevice.”
6. Stanford Artificial Intelligence Laboratory et al., “Robotic operating system.”
7. Harris, C. R., Millman, K. J., Van Der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., ... & Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, 585(7825), 357-362.

Copyright: ©2024 Nicolas Shu, et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.