

PROVIDING COMMON I/O CLOCK FOR WIRELESS DISTRIBUTED PLATFORMS

D. Budnikov, I. Chikalov, S. Egorychev

Intel Labs, Intel Corporation
30 Turgenev St, Nizhny Novgorod, Russia

I. Kozintsev, R. Lienhart

Intel Labs, Intel Corporation
2200 Mission College Blvd, Santa Clara

ABSTRACT

We propose a novel synchronization scheme for distributed audio-video input and output on heterogeneous general purpose computing (GPCs) such as laptops, tablets, PDAs, smart phones, audio recorders, and camcorders. These devices typically possess sensors such as microphones and possibly cameras, and actuators such as loudspeakers and displays. In order to combine them wirelessly into a distributed array signal processing system, it is necessary to provide relative time synchronization to sensors and actuators. In this work we propose a setup and an algorithm to synchronize input and output for a network of distributed multi-channel audio sensors and actuators connected to GPCs. IEEE 802.11 wireless network is used to deliver the global clock to distributed GPCs, while the interrupt mechanism is employed to distribute the clock between I/O devices. Experimental results demonstrate a precision in A/D D/A synchronization precision better than $50 \mu\text{s}$ (a couple of samples at 48 kHz).

1. INTRODUCTION

Arrays of audio/video sensors and actuators such as microphones, cameras, loudspeakers and displays along with array processing algorithms offer a rich set of new features for emerging applications. Until now, array processing required expensive dedicated multi-channel I/O cards as well as expensive high-throughput computing systems due to the requirement to process all channels on a single machine. Recent advances in mobile computing and communication technologies, however, suggest a novel and very attractive platform for implementing these algorithms. Students in classrooms and co-workers at meetings are nowadays accompanied by several mobile computing and communication devices with audio and video I/O capabilities onboard such as laptops, PDA's, and tablets. In addition, high-speed wireless network connections, like IEEE 802.11a/b/g, are available to network those devices. Such ad-hoc sensor/actuator networks can enable emerging applications that include multi-stream audio and video, smart audio/video conference rooms, meeting recordings, automatic lecture summarization, hands-free voice communication, speech enhancement and object localization. No dedicated infrastructure in terms of the sensors, actuators, multi-channel interface cards and computing power is required. Multiple GPCs along with their sensors and actuators co-operate on providing transparent synchronized I/O. However, there are several important technical and theoretical problems to be addressed before the idea of using those devices for array DSP algorithms can materialize in real-life applications. One of the most important problems is to provide a common reference time to a network of distributed computers and their I/O channels.

To illustrate the importance of time synchronization we implemented a Blind Source Separation (BSS) algorithm published

in [1]. In the simplest setting two sound sources are separated using the input of two microphones, each connected to a different laptop. However, without synchronization of A/Ds the BSS algorithm failed to perform separation. Figure 1 demonstrates how a difference of only a few Hz in audio sampling frequency between two channels (laptops) impacts source separation. On the x-axis the sampling difference in Hz between two audio channels at about 16 kHz is shown against the achieved signal separation gain by BSS on the y-axis. As can be seen in Figure 1, a difference of only 2 Hz at 16 kHz reduces the signal separation gain from 8.5 dB to about 2 dB only. In real life the difference in sampling frequency can be even higher as we illustrate in Table 1. BSS is not the only algorithm that is extremely sensitive to sampling synchronization. Other applications that require similar precision of time synchronization between channels are acoustic beamforming and 3D audio rendering.

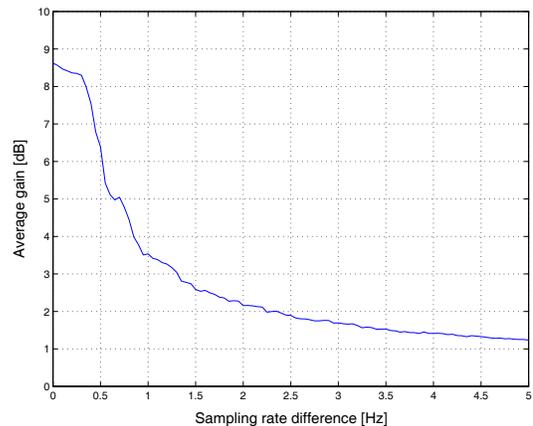


Fig. 1. Sensitivity of acoustic source separation performance to small sample rate differences. Channel 1 is assumed to sample at 16 kHz, while channel 2 is assumed to sample at $16000+x$ Hz. Signal separation gain is calculated for the Blind Source Separation algorithm in [1].

The problem of time synchronization in distributed computing systems has been discussed extensively in the literature in the context of maintaining clock synchrony throughout large geographic areas. Each process exchanges messages with its peers to determine a common clock. Seminal works have been reported in [2] and [3]. However, the results provided there can not be applied directly to our problem, since the precision of time synchronization is too low. NTP, the Network Time Protocol, currently used worldwide for clock synchronization in the best case achieves syn-

| Laptop | Dell Inspiron 7000 | IBM ThinkPad T20 | IBM ThinkPad 600E | IBM ThinkPad T23 |
|-------------------|--------------------|------------------|-------------------|------------------|
| Sampling rate, Hz | 16001.7 | 16003.6 | 16001.8 | 16009.5 |

Table 1. Audio sampling rates of several laptops.

chronization in the range of milliseconds - 2 to 3 orders of magnitudes higher than the microsecond resolution needed for our application scenarios. The Global Positioning system (GPS) provides a much higher clock resolution. Its reported time is steered to stay always within one microsecond of UTC (Coordinated Universal Time). GPS, however, only works reliably outdoors and thus does not completely fit our application scenario. There is also some recent work on synchronization in wireless sensor networks. In [4, 5], the reference-broadcast synchronization method is introduced. In this scheme, nodes send reference beacons to their neighbors based on a physical broadcast medium. All nodes record the local time at which they receive the broadcasts (e.g., by using the RDTSC instruction of the Pentium[®] processor family; the Read-Time Stamp Counter counts clocktics since the processor was started). Based on the exchange of this information, nodes can translate each other's clock. Although promising, the worst case performance of 150 μ s reported in [4] is too high for our application scenario. Our system is similar in spirit but we rely on additional processing to reduce errors in estimation of synchronization parameters.

In general, all clock synchronization algorithms studied in the literature only address the problem of providing a common clock on distributed computing platforms. They do not address specifically for audio) how the I/O can be synchronized with the common clock (we proposed one solution in [6]. In other words, even under the assumption of a perfect clock on each platform, there is still a mechanism required to link the common clock to the data in the I/O channels. On a GPC this is a challenge in itself and we address this problem in this paper.

2. PROBLEM FORMULATION AND SOLUTION

We tackle the problem of distributed I/O synchronization in two steps: (1)(inter-platform) the local CPU clocks of the GPCs are synchronized against a global clock, and (2) (intra-platform) I/O is synchronized against the local clocks and thus also against the global clock. In the experimental results, one of the CPU clocks will arbitrarily be chosen as the global clock.

2.1. Problem Formulation

Each GPC has a local CPU clock (e.g., RDTSC). Let $t_i(t)$ denote the value of this clock on the i -th GPC at some global time t . Assuming a linear model between the global clock and the local platform clock, we get

$$t_i(t) = a_i(t)t + b_i(t), \quad (1)$$

where $a_i(t)$ and $b_i(t)$ are timing model parameters for the i -th GPC. The dependency of the model parameters on global time t approximates instabilities in the clock frequency due to temperature variations and other factors. In practice, these instabilities are

in the order of 10^{-5} . In the rest of the paper we will omit explicit time dependency to simplify our notations. Similarly, the sampling times of audio A/Ds and D/As on GPC's are approximated as:

$$\tau_i(t_i) = \alpha_i(t_i)t_i + \beta_i(t_i). \quad (2)$$

In this model τ_i is simply the number of samples produced by A/D (or consumed by D/A) converter since the start of the audio I/O. Note that two different timing models are required since the audio I/O devices on a typical PC platform have their own internal clock that is not synchronized to other platform clocks such as the RDTSC.

Given the two timing models above the problem that we address in this paper can be formulated as finding $t(\tau_i)$ - the global time stamp of audio sample τ_i . We separate it into two subproblems: finding $\hat{\alpha}_i$ and $\hat{\beta}_i$ such that $t_i(\tau_i) = \hat{\alpha}_i\tau_i + \hat{\beta}_i$ (convert sample number to local time stamp with $\hat{\alpha} = \alpha^{-1}$ and $\hat{\beta} = -\beta/\alpha$) and finding \hat{a} and \hat{b} such that $t(t_i) = \hat{a}t_i + \hat{b}$ (convert value of local clock to global time with $\hat{a} = a^{-1}$ and $\hat{b} = -b/a$).

2.2. Timing relationships on GPC platform

In order to understand the inter and intra platform synchronization methods proposed in this work we briefly describe the operations and timing relationships on a typical GPC. Figure 2 shows a pro-

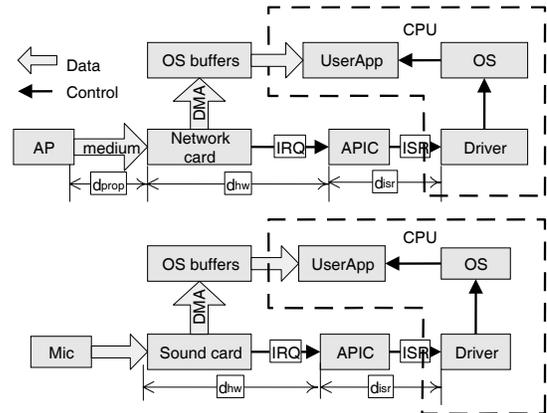


Fig. 2. Network (top part) and audio (bottom part) data and control flows on a typical GPC platform.

cessing diagrams of networking and audio I/O. Both I/O operations have a very similar structure that can be described by the following sequence of actions (only input path is described):

1. Incoming data is received and processed by a hardware device, and eventually is put into a Direct Memory Access (DMA) buffer. This is modeled in Figure 2 by the delay d_{hw} , which is approximately constant for similar hardware.
2. The DMA controller transfers the data to a memory block allocated by the system and signals this event to the CPU by an Interrupt ReQuest (IRQ). This stage introduces variable delay due to memory bus arbitration between different agents (i.e., CPU, graphics adapter, other DMA's).

3. The interrupt controller (APIC) queues the interrupt and schedules a time slot for handling. Because APIC is handling requests from multiple I/O devices this stage introduces variable delay with standard deviation of around 6 *ms* and the maximum deviation of 30 *ms*. Both previous stages are modeled by d_{isr} in Figure 2.
4. The Interrupt Service Routine (ISR) of the device driver is called, and the driver sends notification to the Operating System (OS).
5. The OS delivers a notification and data to the user application(s). This stage has to be executed in a multitasking software environment and this leads to significant variable delays that depend on CPU utilization and many other factors.

In summary, data traverses multiple hardware and software stages in order to travel from an I/O device to the CPU and back. The delay introduced by the various stages is highly variable making the problem of providing a global clock to the GPCs and distributing it to I/O devices very challenging. It is advantageous to perform synchronization as close to hardware as possible, therefore our solution is implemented at the driver level (during ISR) thus avoiding additional errors due to OS processing.

2.3. Providing global clock (inter-platform synchronization)

For the synchronization of CPU clocks over a wireless network we propose to use a series of arrival times of multicast packets sent by the wireless access point (AP). In our current approach we implement a pairwise time synchronization with one node chosen as the master (say $t(t_0) = t_0$). All other nodes (clients) are required to synchronize their clocks to the master¹. A similar approach was also suggested in [4, 5]. Our solution, however, extends it by introducing additional constraints on the timing model. In order to provide a global clock to distributed platforms that is useful to several applications (e.g. joint stream processing and distributed computations) we impose the clock monotonicity condition to make sure that the global clock is monotonically increasing during model parameter adaptation. In addition we smooth the clock model (a_i and b_i in equation (1)) variation by limiting the magnitude of its updates.

The algorithm consists of the following steps:

1. AP sends next beacon packet.
2. Master node records its local time of packet arrival and distributes it to all other nodes.
3. Client nodes record both their local times of arrival of beacon packets from AP, and the corresponding times received from the master.
4. Clients update local timing models based on the set of local timestamps and corresponding master timestamps.

Let us assume that in Figure 2 the packet j arrives to multiple platforms approximately at the same global time corresponding to local clocks t_i^j ($d_{prop} \approx 0$). The set of observations available on the platforms consist of pairs of timestamps $(\tilde{t}_0^j, \tilde{t}_i^j)$. From Figure 2 we have $\tilde{t}^j = t^j + d_{hw} + d_{isr}$ (we omitted dependency on i) that we further model as $\tilde{t}^j = t^j + d + n$. In this approximation d models all constant delay component and n represents the

¹A more complex approach of performing joint timing synchronization is potentially more accurate.

stochastic component. Given the set of observations $(\tilde{t}_0^j, \tilde{t}_i^j)$ we are required to estimate the timing model parameters \hat{a}_i and \hat{b}_i for all slave platforms. In our experiments a window of 3 minutes is used to estimate current values of \hat{a}_i and \hat{b}_i using the least trimmed squares (LTS) regression [7]. LTS is equivalent to performing least squares fit, trimming the observations that correspond to the largest residuals (defined as the distance of the observed value to the linear fit), and then computing a least squares regression model for the remaining observations. Figure 3 shows comparison of quantiles of residuals with quantiles of normal distribution and Figure 4 plots the histogram of residuals. The distribution appears to be close to Gaussian except for the presence of a few outliers (see Figure 3) that do not fit into a normal distribution. The trimming step is specifically targeted to remove those outliers.

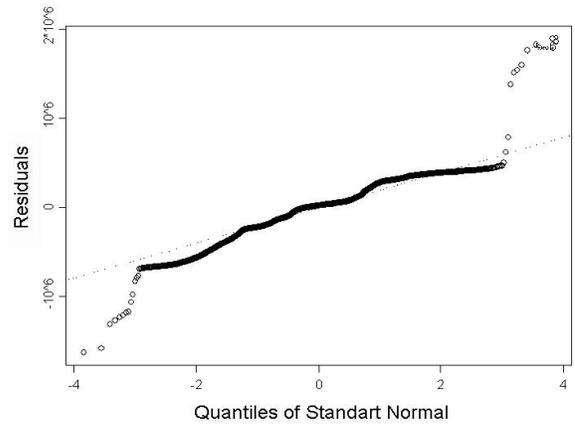


Fig. 3. Comparison quantiles of residuals with quantiles of the normal distribution. Points away from the straight line are treated as outliers and removed during regression.

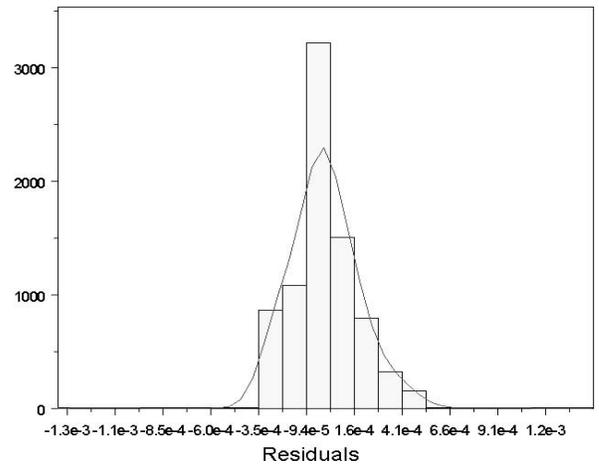


Fig. 4. Histogram of residuals and the normal probability density function.

2.4. Synchronizing audio to CPU clock (intra-platform synchronization)

In order to synchronize the audio clock to the CPU clock we use a similar approach as the one presented in the previous section. The ISR of the audio driver is modified to timestamp the samples in the OS buffer using the CPU clock to form a set of observation pairs $(\tilde{t}_i^j, \tau_i^j)$, where j now represents the index of an audio data packet. Following our model in Figure 2 we have $\tilde{t}^j = t^j + d_{hw} + d_{isr}$ (we omitted dependency on i) that we further represent as $\tilde{t}^j = t^j + d + n$. Except for the fact that the τ^j are available without any noise (it is simply the number of samples processed!) we are back to the problem of determining the linear fit parameters for pairs of observations that we solved in the previous section using the LTS method.

In summary, by using LTS procedure twice both local and global synchronization problems are solved and the audio samples can be precisely synchronized on the distributed GPCs.

3. EXPERIMENTAL RESULTS

The distributed test system was implemented with several off-the-shelf laptops implementing Intel®Centrino™ Mobile Technology using the following software components (see also Figure 5): (a) A modified WLAN card driver timestamps each interrupt, parses incoming packets in order to find all master beacon frames, and stores their timestamp values in a cyclic shared memory buffer. The timestamp values as well as the corresponding message IDs are further accessible through the standard driver I/O interface. (b) A modified AC97 driver timestamps ISRs and calculates the number of samples transmitted since the beginning of the audio capture/playback. The value pair is placed into a cyclic shared memory buffer. (c) The *synchronization agents* are responsible for synchronizing the distributed system. We have three types of agents: the multicast server (MCS), the master synchronization agent (SAM) and the slave synchronization agent (SAS). The MCS periodically broadcasts beacon packets (short packets with unique ID as the payload). The SAM and SASs use the modified WLAN driver to detect the beacons. The SAM periodically broadcasts its recorded timestamps of beacon arrivals to the SAS devices. Based on SASs' recorded timestamps and the corresponding SAM timestamps, each SAS calculates the clock parameter to convert between the platform clock and the global clock. The clock parameters are placed in shared memory for use by other applications. (d) The *Synchronization API* allows user applications to retrieve the local clock value, access the clock parameters, and convert between the platform and global clock. (e) The *audio API* allows user applications to retrieve pairs of local timestamps and sample numbers, as well as to convert global timestamp values to sample numbers and vice versa. It also provides transparent synchronized capture and playback.

Based on these components a distributed audio rendering system was implemented with three laptops (see Figure 5). The first laptop was used as the MCS. Modified AC97 and WLAN drivers were installed on other two laptops. SAM was started on the second laptop, while SAS were started on the third laptop. The distributed system was instructed through the audio API to synchronously playback a Maximum Length Sequence (MLS) signal on the two synchronized laptops. The line-out signals of both laptops were recorded by a multichannel soundcard. The measured inter-GPC offset was at most 2 samples at 48 kHz ($\leq 42 \mu\text{s}$).

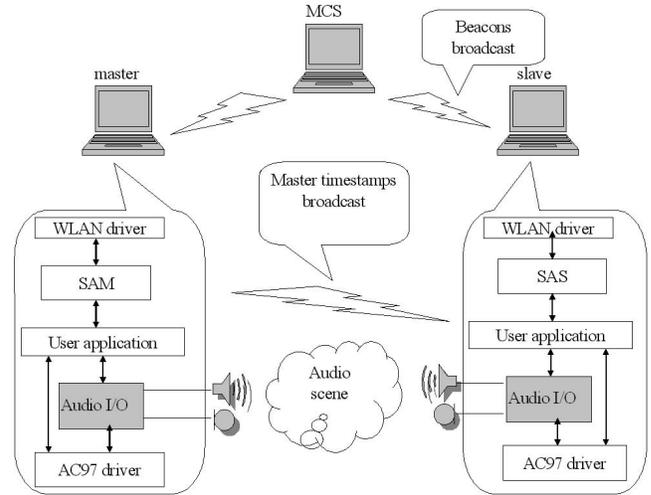


Fig. 5. Distributed audio rendering/capturing system setup

4. SUMMARY

We have proposed a general two-step solution to provide a common I/O clock for distributed platforms. In the first step, the CPU clocks of the distributed GPCs are synchronized against each other by using the wireless network. In the second step, each GPC is modeled as a distributed system with multiple I/O devices interconnected to the main memory and CPU via shared buses and a similar synchronization principle is applied. Our results show that at any time the audio I/O offset between different GPCs can be kept below $50 \mu\text{s}$ that enables usage of *distributed GPCs* for array signal processing of audio and video.

5. REFERENCES

- [1] C. Fancourt and L. Parra, "The coherence function in blind source separation of convolutive mixtures of non-stationary signals," in *Proc IEEE Workshop on Neural Networks for Signal Processing*, 2001, pp. 303–312.
- [2] L. Lamport and P.M. Melliar-Smith, "Synchronizing clocks in the presence of faults," *JACM*, vol. 32, no. 1, pp. 52–78, 1985.
- [3] D. Mills, "Internet time synchronization: the network time protocol," *IEEE Tran Comm*, vol. 39, no. 10, pp. 1482–1493, 1991.
- [4] M. Mock, R. Frings, E. Nett, and S. Trikaliotis, "Clock synchronization for wireless local area networks," in *IEEE 12th Euromicro Conference on Real-Time Systems (Euromicro RTS 2000)*, 2000, pp. 183–189.
- [5] J. Elson, L. Girod, and D. Estrin, "Fine-grained network time synchronization using reference broadcasts," in *5th Symposium on OS Design and Implementation*, Dec 2002.
- [6] R. Lienhart, I. Kozintsev, and S. Wehr, "Universal synchronization scheme for distributed audio-video capture on heterogeneous computing platforms," in *Proc 11th ACM Conf on Multimedia*, 2003, pp. 263–266.
- [7] P. J. Rousseeuw, "Least median-of-squares regression," *JACM*, vol. 79, pp. 871–880, 1984.