PhonePeti: Exploring the Role of an Answering Machine System in a Community Radio Station in India

Zahir Koradia
Dept. of Computer Science and Engineering
IIT Bombay
Mumbai, India

Aaditeshwar Seth
Dept. of Computer Science and Engineering
IIT Delhi
New Delhi, India

ABSTRACT

Community Radio (CR) stations are short range radio stations that serve the local media needs of their surrounding communities. Community participation by way of helping set the station agenda, airing of people's voices, and providing them with a local communication medium, is the defining feature of CR. But this philosophy has been hard to execute in practice because of logistical difficulties, with station staff not being able to reach out to a listenership-base spread across several hundreds of square kilometers. In today's context though, the high penetration of mobile phones has made it easier for listeners to participate in the running of radio stations, but the potential of telephony and radio integration has been exploited only minimally.

In this paper, we explore the use of PhonePeti, an automated answering machine system in a community radio station based in Gurgaon, India. Answering machines are one of several ways to bring together the radio and telephony mediums. We show that this alone has the potential to considerably improve community engagement, but it also opens up many interesting issues on usability. Through quantitative and content analysis of 758 calls from 411 callers over two iterations of PhonePeti, combined with telephonic interviews of several callers, we show that significant challenges arise in being able to explain the concept of an answering machine to people who have not been exposed to a similar system in the past. We then show, through call statistics, that PhonePeti has increased community engagement by enabling more listeners to reach the station. Finally, we show that an answering machine system can be used to collect useful information from the callers.

1. INTRODUCTION

Community Radio (CR) stations focus on meeting the local media needs of their surrounding communities. Generally operating on the FM band in most countries, these stations have a small reach of 10-15 km radius, and air locally created programs related to folk music, civic issues, careers, health and hygiene, etc. While actual organizational structure of CR stations may vary around the world, community participation is considered a core principle in CR station operations.

CR stations strive to enable their communities to establish the station agenda, participate in daily activities such as content creation, editing and broadcast, and feel empowered to voice their opinions and demand. Such community participation helps ensure that a station's content reflects the needs of the community in which it is located. In addition, as Sterling et al. point out, participation of under-represented sections of the community in content creation leads to their social development [14].

The wide proliferation of mobile phones has made it easier for CR stations to engage with their listeners. Listeners can now simply call the station and express their opinions, something they could otherwise do only when they met a station staff member in person. Prior work exists on CR being able to utilize the growing penetration of mobile phones [9, 3], but we believe that the potential of radio and telephony integration has not been fully explored. The particular gap on which we focus is the role an interactive voice response (IVR) system can play in a CR station, and how the community at large can be trained to use such a system.

In this paper, we present PhonePeti, an answering machine system deployed at a CR station located in north India, called Gurgaon Ki Aawaz (GKA). In the study of PhonePeti usage, we attempt to answer the following questions:

1. What kind of messages do the callers leave? How often do they call? Does PhonePeti really improve community engagement?

2. Are callers able to understand that PhonePeti is an automated system? How do they learn to use it? Are the usage instructions aired on radio helpful at all?

3. Can an answering machine system be used to solicit information from the callers and conduct brief surveys? How many callers answer the survey questions?

We present an analysis of the PhonePeti deployment at GKA, studied over a period of five months. We show how PhonePeti has become a standard way for community members to reach the station, particularly during off-office hours when no staff is available to answer the regular office phone. Through an analysis of 758 phone calls received over two iterations of PhonePeti, we describe the types of messages that community members leave on the system. We also show

that being able to leave a message to an automated voice does not come naturally to the callers – more than 60% of the calls did not contain any useful audio, and more than 50% of the callers called into the system only once. Finally, we show how PhonePeti can be used to solicit structured feedback from the radio station listeners.

The rest of this paper is organized as follows. In the next section we present a survey of related work, followed by a brief introduction to Gurgaon Ki Aawaz in Section 3. In Section 4, we present the PhonePeti design and describe results obtained in the first iteration of its deployment. Section 5 presents the second iteration of PhonePeti followed by a summary of the results in Section 6. We then present the limitations of our work in Section 7, highlight the broader scope for radio-telephony integration in Section 8, and finally present the conclusions of our study in Section 9.

2. RELATED WORK

Related work can be categorized into (a) the use of IVR systems in developing regions, and (b) use of telephony in community radio stations.

2.1 IVR systems in developing regions

Several recent studies explore the use of IVR systems in developing regions as a way of information creation and dissemination. Agarwal et al. [5] discuss a VoiKiosk system where callers could set up audio blogs and create advertisement for their services over a simple phone call. Patel et al. [12] describe the evolution of a voice based question and answer forum into a social networking platform in Gujarat, India. Lerer et al. [10] explore the challenges of conducting voice based surveys without giving prior training to the participants on its usage. Sherwani et al. [13] and Grover et al. [8] study IVR systems for providing health related information, and Medhi et al. [11] study voice user interfaces for banking activities. Our work is similar, with the primary difference being the community radio context of its deployment.

2.2 Use of telephony in community radio stations

The increased penetration of mobile phones in developing regions has encouraged CR stations to consider telephony as an opportunity to increase community participation. The GRINS [9] radio automation system allows CR stations to make and receive calls, screen incoming calls from listeners, and put calls live on air. Several CR stations in India have used GRINS in multiple ways, for example, to obtain live feedback from listeners, put civic authorities on air, and record audio snippets by famous personalities. GRINS is able to record conversations between station staff and a caller, and optionally put the conversation live on air. Our work is complementary – we study the use of offline recording systems where one end of the call is handled by a computer instead of a station staff member.

FreedomFone [3] allows CR stations to build simple IVR systems, conduct polls, and allow callers to leave messages for the station. A good example is Farm Radio International's AFRRI project [1] in Tanzania, where FreedomFone helped make archived radio programs available to callers over the phone, and an answering machine capability was used to collect feedback about the radio programs. However, only an informal report of the study is available [15] with many experiment details missing. In comparison, we present detailed system usage and call analysis.

3. GURGAON KI AAWAZ

Gurgaon Ki Aawaz (GKA) is a community radio station located in Gurgaon, India. The primary listener base of the station are migrant workers employed in Gurgaon. GKA airs a large variety of programs on social issues, micro-entrepreneurship, careers, and financial instruments for the poor. Most of the programs on GKA are aired in Hindi and *Haryanvi*, a local dialect of Hindi. Local folk songs, called *Ragini* are extremely popular; the station airs a two hour program daily consisting only of *Ragini* requests. True to the spirit of community radio, the station has always hired staff from the community around the station, and trained them in content creation, editing, and broadcast.

3.1 Community Participation

GKA enables community participation through two main activities: (a) field recordings and (b) listener phone calls. More than 80% of the content for programs is collected from the field, where the station staff record stories from the community and collect feedback on the existing programs. In addition, listeners call the station to make song requests, obtain additional information about programs, alert the station staff of technical broadcasting problems, or even complain about civic facilities. The station staff record many of these phone calls via a hand held recorder and put the recordings on air.

According to the log books maintained by the station, GKA had received more than 8000 calls till February 2011, giving an average of more than 20 calls per day. The station received an average of 30 calls per day between December 2010 and February 2011, indicating that the number of calls per day has increase since the station started its operations.

With just one phone line available at the station though, callers have often complained of the phone line being busy. In addition, the staff is available at the station to receive phone calls only during office hours. This limitation makes it difficult for listeners such as taxi drivers (who work late hours) to call into the station.

To address these issues, we deployed PhonePeti at Gurgaon Ki Aawaz. PhonePeti was designed enable more listeners to leave messages for the station and allow listeners to reach the station 24 hours a day.

For the purpose of this paper, we define *community engagement* as the number of calls received per day by the station from the listeners. Although this definition takes a very narrow view of community participation by not considering qualitative aspects of phone conversations and completely ignoring other forms of participation, we choose this definition as it is easy to quantitatively measure the impact of PhonePeti on *community engagement*. For example, *community engagement* at GKA til February 2011 was 20, and that between December 2010 and February 2011 was 30.

We next describe the first iteration of PhonePeti deployment.

4. PHONEPETI 1

The first iteration of PhonePeti (*P1*) was deployed for a period of four months from 6 January 2011 to 11 May 2011. Below, we describe research questions of interest, the overall

architecture, experiment design, and results.

4.1 Research questions

We wanted to understand the following aspects about *P1*'s usage, which are unique because of the community radio context in which the IVR system is deployed.

1. What times of day do the listeners call? How many calls are made? What is the content of the recorded messages? Call volume would help us understand the impact of *P1* in improving *community engagement* and content analysis would allow us to gain an understanding of PhonePeti usage.

2. Do the callers understand how to use the system? How do they learn? We wanted to confirm whether listeners can learn how to use PhonePeti just by hearing the information aired on radio about the system. This is likely to be different from typical IVR deployments where users are trained in-person to use the system, or learn through word of mouth from other users, or simply through trial and error.

3. What is the impact of gender and familiarity of prompt voice on the recorded messages? This is interesting to explore because the listeners are already used to hearing regular station staff on air; if they heard the prompt in a *familiar* voice, it may reduce their inhibitions about the system and they may speak more freely.

We next present the architecture of the PhonePeti deployment at GKA.

4.2 Architecture

To setup PhonePeti at GKA, we deployed a simple IVR system at our premises, as shown in Figure 1. PhonePeti is implemented as an application over Asterisk [2], a telephony engine that can be programmed to respond automatically to the caller by presenting IVR instructions. A GSM gateway is connected to the computer running Asterisk to terminate GSM calls and redirect them to Asterisk for further processing. Thus, when a listener calls the mobile number associated with PhonePeti, the call lands on Asterisk via the GSM gateway. At this point, Asterisk begins execution of the telephony application which prompts the caller to leave a message after a beep. The application records the message from the caller and pushes it to a web server to make it available to the station staff over the Internet. Loudblog, an audio blogging and pod-casting software [4], is used to list the recordings in an easily browsable manner, and allow the staff to categorize and selectively publish the messages on their own website. A few selected recordings published by GKA are available at `http://bit.ly/phonepeti`[1]. The station staff can also download the recordings on their local machine for broadcasting.

We next present the experiment design to help us answer the research questions of interest.

4.3 Experiment design

The experimental setup for PhonePeti was designed to (a) capture call related statistics, (b) record messages left by callers, (c) provide tools for the analysis of content, and (d) allow the changing of prompts for different callers.

[1]A shortened URL used to obscure identification information for blind review

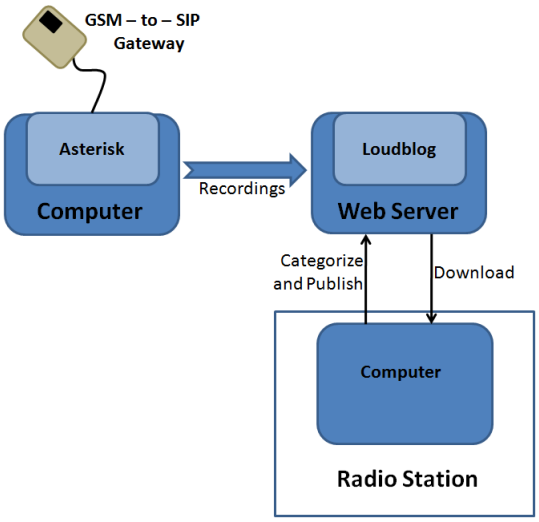

Figure 1: PhonePeti Design

4.3.1 Information collection and processing

To conduct a usage analysis of *P1*, we wanted to know the number of calls received from a caller, the times of day at which they called, and the recorded messages. The *P1* telephony application therefore logged this information in a database.

Once the recorded messages were pushed to the web server, we categorized them based on their content. Categories included *song requests, suggestions, contributions for programs, etc*, the list having been finalized after scanning a few initial recordings. The content categorization exercise was done in two passes. In the first pass, the station staff categorized the messages after an initial training session. In the second pass, we selected a few recordings randomly and double-checked the categorization to ensure accuracy and coder reliability.

To gain insights into the usage patterns we observed, we even conducted telephonic interviews of several callers. This included both callers who had called only once, and those who had called several times. While we intended to conduct semi-structured interviews, doing so turned out to be extremely hard, especially for one-time callers. Many callers did not pick-up our calls, and several callers refused to acknowledge that they had even called the system. This could be because their friends or relatives may have made the call, or because they did not feel comfortable in revealing information to us. Multiple time callers, on the other hand, were more comfortable in talking to us.

4.3.2 Call flow

Figure 2 shows the flowchart of a call in *P1*. When a listener calls into the system, a prompt is played asking the caller to leave a message after a beep. If the caller had called in the past then the same prompt is selected each time for that caller; else, a prompt is chosen at random from a pool of prompts. Details about the pool of prompts is presented later in the section.

The caller could end the recording of his message by pressing a button between 0 and 9. Alternately, if silence was detected for a period of 3 seconds, or the caller had recorded audio for the maximum allowed duration of 1 minute, then the recording was automatically terminated. The caller's recording was then played back to him[2], followed by an acknowledgement message that the recording had been archived successfully. Finally, the system saved the call information in a database, pushed the recorded audio to the web server, and terminated the call.

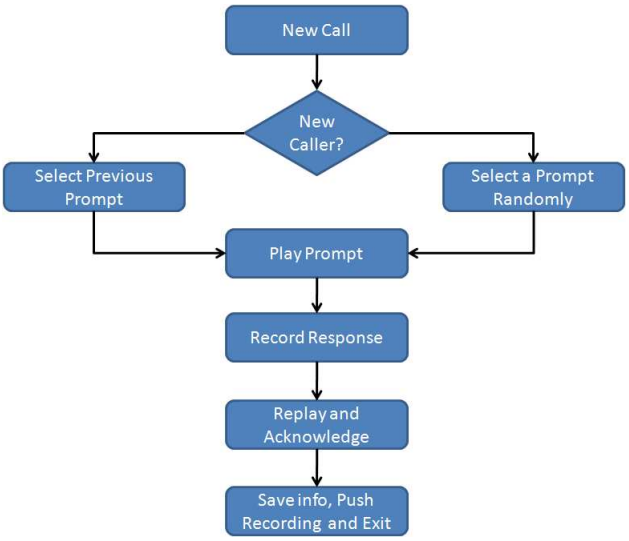

Figure 2: Call flow of PhonePeti 1 (P1)

4.3.3 Pool of prompts

To observe the impact of gender and familiarity of prompter voice on the messages recorded by the callers, we recorded prompts from four people: (a) Familiar-Female, (b) Familiar-Male, (c) Unfamiliar-Female, and (d) Unfamiliar-Male. Familiar prompters were those of GKA staff who regularly spoke on air. The listeners had never heard the unfamiliar prompters before. These categories are shown in Table 1.

Table 1: Prompters used in *P1* based on familiarity and gender.

	Familiarity	
Gender	Familiar-Female	Unfamiliar-Female
	Familiar-Male	Unfamiliar-Male

The prompt recorded by each prompter is transcribed below. Whenever a new caller called into *P1*, one of these four prompts was chosen at random and attached to the caller.

> *Welcome to Gurgaon Ki Aawaz. If you want to ask a question or leave a suggestion, then record your message after the beep.*

[2]We consciously use masculine gender for callers since all but one caller in PhonePeti were males.

4.3.4 Introduction of the system to listeners

Since we wanted to study how comfortable callers felt in interacting with a machine rather than a human operator at the other end, PhonePeti was presented to the listeners as another way for them to reach the station. A short radio advertisement scripted as a drama was used to tell listeners about PhonePeti: a "number" where listeners could record complaints, suggestions, song requests, and questions for the station. The drama also highlighted the system as being available 24 hours and being machine operated. The transcript of the advertisement is presented below.

Night time indicated by crickets making sounds

Wife: Close the door. Can't you see the sewage canal is open. It smells so awful, one can't even sleep here.
Husband: Then record a complaint against the municipal corporation at Gurgaon Ki Aawaz.
Wife: But they do it only until 5 in the evening.
Husband: No no. That is not true. Now questions will be recorded 24 hours. Not only questions, but also suggestions, and song requests too!
Wife: So the number must be the same?
Husband: There are two numbers. To talk to a machine dial .....

Audio conveying that the wife dials a number, hears the prompt, and records the complaint

Husband: After recording, to listen to what you have recorded you will have to press one more button.

Audio conveying that a button is pressed, followed by a repeat of the wife's recorded complaint, and an acknowledgement by the system.

Wife: Oh wow! Distances have indeed reduced because now our voice is Gurgaon's voice, 24 hours!

GKA airs a live program "Gurgaon Ke Haal" every morning, where the station staff discuss locally relevant issues from newspapers, and ask listeners to call in and record their views about the issue. The PhonePeti number was sometimes advertised on this program. It was also advertised a few other times in different contexts, such as alongside a two hour program on requests for folk songs, but the exact frequency and advertising context was controlled by the station staff and not by us. This lends a need-based character to PhonePeti's usage: it was the station staff who recognized its use in soliciting feedback from the listeners, and independently decided when an automated system would be useful as compared to a human operator conversing with callers.

Occasionally listeners also called the office phone to understand the PhonePeti service. Some would specifically ask "Whom will I talk to on that number?", and the station staff would reply that it was a machine that will receive the call.

Thus, PhonePeti was presented as an alternate way for the listeners to reach the station.

4.3.5 Utilization of Recordings

Recordings from PhonePeti were used by the station staff in the same manner as the recordings from the office phone. GKA runs a feedback program that contains messages left by the callers for the station. Before PhonePeti was set up, the program contained messages recorded on the office phone only. After PhonePeti deployment, it was changed to additionally contain messages left by callers on PhonePeti.

We ensured that the feedback program made no distinction between the recordings obtained from PhonePeti versus those obtained from the office phone. We did not control the the ratio of PhonePeti recordings to office phone recordings in the program as we felt that such constraints would impact the quality of the program content.

We next present the results of *P1*'s usage.

4.4 Results

P1 was deployed for over four months between 7 January 2011 and 11 May, during which 320 calls were received by the system from 169 callers[3].

4.4.1 High level analysis

Our analysis of the log files and the database showed that many calls were terminated even before the recording began. We term these calls as *null* calls. Analysis of the recorded messages revealed that several callers did not record anything; we termed these as *empty* calls containing only background noise. Yet another set of recordings contained callers repeating "Hello...Hello" as if trying to elicit a response from the other end. We term these calls as *hello* calls. Collectively we categorize *null*, *empty*, and *hello* calls as *bad* calls, and all the other calls containing valid recordings as *good* calls.

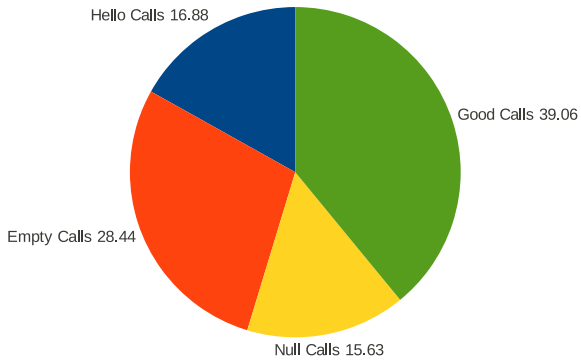

Figure 3: Percentage of Good and Bad calls in *P1*. A total of 320 call were received.

Figure 3 shows the distribution of *good* and *bad* calls in *P1*. *Bad* calls account for more than 60% of the total calls made. We therefore wanted to understand the reasons that could have led to *bad* calls, and attempt to address the issues.

We first looked at the causes for *null* and *empty* calls by conducting telephonic interviews of the callers. We attempted to talk to more than 20 callers, but were able to get responses from only 9 callers. A variety of reasons were cited: (a) the caller did not understand the system, (b) technical failure at the caller end, (b) the caller dialed just to check how the system sounded, (c) some other work came up for the caller during the call, (d) the caller got scared and confused in dealing with an automated system, and (e) the caller was hesitant to talk to a radio station. While interviewing these callers, we got the impression that a lack of information about who was present at the other end of the call, and how the recording would reach the staff, were a source of considerable anxiety. We realized that the radio advertisement could have been modified to explicitly clarify these aspects. We made these modifications in the second iteration of PhonePeti, described in Section 5.

The case of *hello* calls was more easily explainable – the caller was clearly expecting a person at the other end instead of a computer. One reason for this could have been that the caller was not aware that PhonePeti was an automated system. While the radio advertisement mentioned that PhonePeti was machine driven, the prompt gave no such indication. We realized that the prompt could be modified to reiterate the automated aspect of the service. We tested this aspect in the next iteration of PhonePeti, described later. Another reason for *hello* calls could have been that the callers had no experience of interacting with an IVR system. Unfortunately we could not get a conclusive validation of these reasons through phone interviews with *hello* callers.

4.4.2 Community Engagement

We next look at contribution of *P1* to *community engagement* at GKA. Recall that we defined *community engagement* at a station as the number of calls received per day by the station from its listeners. Also recall that before PhonePeti was deployed, *community engagement* at GKA was 30. In *P1*, 320 calls were received over four months, which translates to 3.1 calls per day, or *community engagement* increment of 10%. We call this *potential* increment, since this would be the increment for GKA if all the calls were *good*.

However, if we consider only *good* calls then the increment reduces to just 3%. We call this *real* increment as it corresponds to actual increase in the interaction between the station and its listeners. Thus, *P1* has made little *real* improvement in *community engagement*, but has the potential to contribute more if *bad* calls can be converted to *good*.

4.4.3 Time of day distribution of calls

We now look at the distribution of calls across the time of day. Figure 4 show that more than 66% of the total calls were made between 6 p.m. and 8 a.m. the next day. This was the time when the station office was closed and there was no one to receive calls on the office phone. This is a qualitative indicator of increased *community engagement*, since it indicates that the increased number of calls is not just a result of providing an extra phone number to the listeners but also a result of the availability of the number 24 hours a day.

It also interesting to see that some calls were also received between midnight and 4 a.m. This was not surprising for the station staff since many of the station's listeners are cab drivers who have night duties ferrying passengers to and from the airport nearby. Finally, a large number of calls were made between 10 a.m. and 12 noon, which is when the folk song request program was aired.

[3]We assume a one-to-one mapping between a phone number and a person.

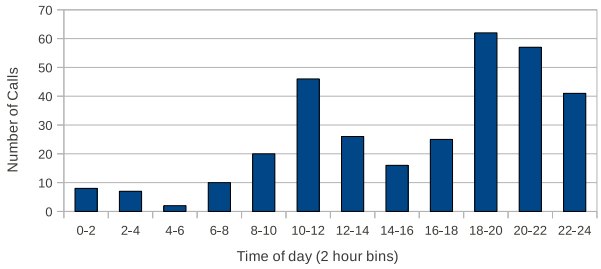

Figure 4: Distribution of calls made in *P1* across time of day.

4.4.4 Distribution of number of calls per caller

Figure 5 shows a CDF of the number of calls per caller. More than 50% of the callers called into *P1* only once, and less than 10% of the callers called four or more times.

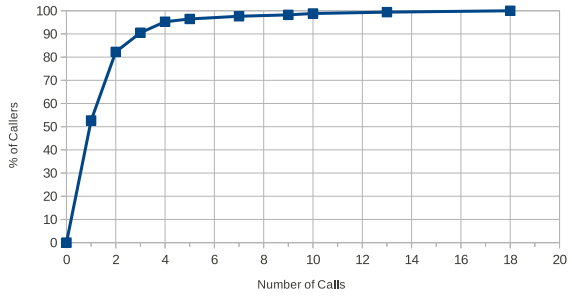

Figure 5: CDF distribution of number of calls made by the callers in *P1*. There were a total of 169 callers.

To understand why so many callers called only once, we again attempted to contact several one-time callers over the phone but failed to get in touch with enough people. The only insight we could gather was that there was some anxiety among one-time callers about how the system worked in conveying their messages to the station staff.

It is interesting to note that Agarwal et al. [5] observed a similar pattern in the distribution of number of advertisement recordings per caller: among 3255 callers, about 52% recorded an advertisement just once. While the authors do not attempt to explain the cause for this behavior, the similarity of their results to ours allows us to build an understanding of call frequencies in IVR systems designed to allow callers to record messages.

4.4.5 Learning to use PhonePeti

We next describe how the callers learned to leave a *good* message in *P1*. We looked at the first three calls from all the callers who had called three or more times into *P1*. There were 30 such callers. Amongst these, we pruned the list to those who had understood the system in the first call itself, that is, those who had left a *good* message in the first call. There were 8 such callers, leaving 22 who had not understood the system in their first call. We then analyzed the second and third calls of these 22 callers to see if they were eventually able to learn to use *P1* and record a *good* message. The division of callers among (a) those who understood the system in the first call, (b) those who learned the system by their third call, and (c) those who learned the system by making more than three calls, is shown in Figure 6. We found that 11 out of 22 were able to learn the system by their third call. However, note that only 30 out of 169, or 18% of the callers, called into the system three or more times. This shows that while learning can happen over time, it is not fast enough, since before a caller is able to learn he often chooses not to call again.

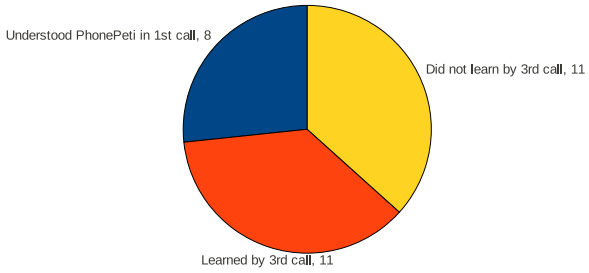

Figure 6: Learning PhonePeti: Three- or more-time callers divided based on number of calls needed to learn using PhonePeti

We next describe how these 11 callers learned to use *P1* by conducting telephonic interviews of four callers. Two reported that the call recording ended automatically in their first call. We suspect that silence detection in the telephony application could have terminated the recording prematurely. The others reported that when they heard the prompt they did not know what to do, implying that they had not understood the system based on the radio advertisement instructions. When asked about how they learned to leave a *good* message, three of the interviewees said that they learned by listening to the radio program again, and one said that the system's similarity to standard IVR systems used in call-centers helped him understand *P1*.

Thus it seems that radio advertisements were the main source of learning. We would like to add though, that in a couple of recordings we could hear a companion in the background trying to guide the caller to use the system. It is hard for us to differentiate between the impact of radio advertisements and that of companion trainers, from our current data.

4.4.6 Content analysis of good calls

A large variety of messages were recorded by the callers, including song requests, appreciation of the station's work, messages for civic authorities, poem or joke contributions, and messages for other listeners. Figure 7 shows the categorization of *good* calls.

The highest number of calls were made for song requests. This was expected since the folk song request program is the most popular. A large number of suggestions were also received, many of which revolved around availability of the

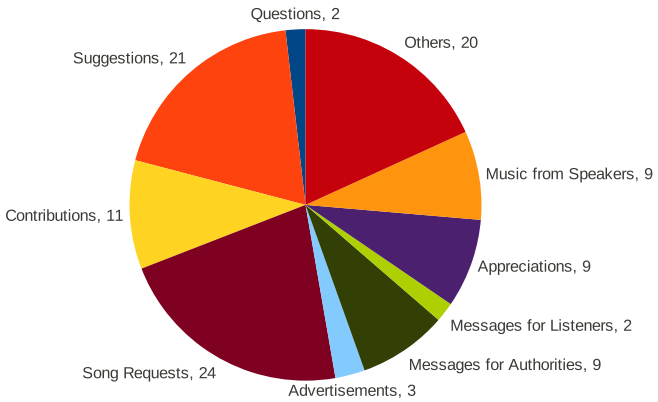

Figure 7: Percentage wise distribution of good calls in *P1*. 110 good calls were received in all.

radio signal in different surrounding suburbs. Some contributions of poetry and jokes were also received. We present below a poetry contribution as a representation of the rapport established between the station and its listeners.

> *This station itself does not know*
> *how much of a support it is for the people,*
> *Not only the world*
> *but it is dearer than our lives to us,*
> *It does not matter what others say,*
> *It was ours yesterday, today also it belongs to us.*
> *May its programs always remain near our ears,*
> *May be special, its relationship with us.*

Several messages were also addressed to civic authorities on water logging and sanitation issues, showing that the listeners were comfortable in using the CR station as a proxy to convey their grievances to the administration. One such message is transcribed below.

> *Through GKA, I want to say that the urban development authority has installed a sewage treatment plant near Chandan Nagar, which throws out waste water using a motor. The pipe connected to this motor has been elevated so high that its smell gets carried in the wind. Living in the surrounding areas has become very difficult. I request you to forward my complain to the authority.*

Thus, a large variety of relevant messages were received though *P1* demonstrating that IVR systems can improve community engagement for the radio station.

Analysis of the *good* calls showed little impact of gender and familiarity of prompt voice on the messages. We had initially suspected that callers may feel intimidated or scared while recording complaints against the authorities or the station, and expected prompt gender or familiarity to have an effect. However, during our analysis of these calls we found no signs of inhibition like non-disclosure of name and location, or hesitation in recording the message. These results are similar to those obtained by Evans and Kortum [7], where disclosure rates were not affected by prompt gender or personality in an IVR system deployed in a medical setting. As a result, we chose not to explore the impact of prompt gender and familiarity in greater detail.

4.4.7 Recording termination

We next looked at a specific aspect of IVR interfaces: termination criteria for recordings. There are four ways of terminating the recording of a message: (a) pressing a key between 0 and 9, (b) silence detection for 3 seconds by the system, (c) call hangup, and (d) maximum record duration (1 minute) exceeded.

When using silence detection, there is a trade-off in choosing the duration of silence. If the duration is too short, the caller may not get enough time to think before recording a message. This problem is more acute in cases where the callers do not have prior experience of using IVR systems. On the other hand, having a long detection duration will make the caller wait that much longer before he hears back his recording and the acknowledgement message. Using a key press as termination criterion avoids these problems, but callers who do not have prior experience of IVR systems may not press a key after recording the message. In *P1*, we allowed both silence detection and a key press as termination criteria, and studied which was used more often.

Table 2: Mode of recording termination for 295 calls that started recording responses.

Mode of recording termination	Call count
DTMF Digit 1	17
DTMF Digit 2	0
DTMF Digit 3	4
DTMF Digit 4	4
DTMF Digit 5	13
DTMF Digit 6	4
DTMF Digit 7	2
DTMF Digit 8	5
DTMF Digit 9	3
DTMF Digit 0	0
Silence Detection	160
Call Hangup	80
Max Duration Exceeded	3

Table 2 shows how the recordings were terminated for 295 calls. Silence detection was the primary mode of termination, accounting for more than 50% calls. This was in spite of instructions given in the radio advertisement to press a button to terminate the recording. The prompts however did not carry this instruction.

As explained earlier, such a high percentage of terminations through silence detection could have had one adverse effect: callers may not have been able to begin recording a message in time, and silence detection would have terminated the recording resulting in *empty* calls. To see if this was indeed the case, we analyzed the empty calls of duration less than or equal to 3 seconds. Out of 76 such calls, 51 were terminated due to silence detection. This probably indicates that 3 seconds of silence detection duration was too aggressive and could be relaxed. This could of course impact those cases when a caller had completed recording his message and was actually waiting for silence to be detected. Therefore, explicitly encouraging callers to press a button to end the recording may be a better approach. We

tested this in the next iteration of PhonePeti, described in the following section.

5. PHONEPETI 2

Usage results of *P1* showed that a large number of calls were *bad*, resulting in acceptable *potential community engagement* but poor *real community engagement*. Many of the calls were *empty* because of aggressive silence detection threshold. Some *hello* calls could also have been avoided by changing the prompt to reiterate the machine-driven nature of the service. We therefore wanted to revise *P1* to improve *real community engagement*. In addition, GKA wanted to collect some information about their listeners, specifically, what kind of ladies comprise their listener base, and what were preferred program schedules of the listeners. We therefore revised the radio advertisement, the prompts, and the call flow, in the second iteration of PhonePeti (*P2*). We focused on two specific aims: (a) improve *real community engagement*, and (b) obtain specific information from the callers as desired by GKA. We next describe how these objectives were incorporated in the experiment design, followed by results.

5.1 Improving community engagement

Recall that we noticed in our *P1* experiments that many *bad* calls happened because the callers did not know who they were talking to at the other end, and how their recordings would reach the station. To address this issue, we expanded the on-air advertisement and highlighted that (a) PhonePeti is a computerized service, (b) when a caller calls, he talks to a machine, and (c) the recorded message becomes available to the station on their studio computer, from where they can broadcast the audio. We increased requested the station to significantly increase the advertisement broadcast frequency to increase the number of *good* calls. Our analysis of broadcast statistics showed that the frequency was increased from 17 per week to 63 per week.

We used only the familiar-female prompt voice to keep the experiment design simple. We modified this prompt to explicitly mention that PhonePeti was a computerized service. And we also modified the prompt to explicitly instruct the caller to press a button between 0 and 9 to terminate the recording. By doing so, we hoped to increase call terminations through key press, thus allowing us to use less aggressive silence detection in the next iteration.

5.2 Collecting specific information

GKA intended to change its program schedule for which it wanted feedback from the callers about their preferences. In addition, GKA was planning to start a program containing interviews of female listeners, and hence needed access to them. To collect this information from the callers, we modified the PhonePeti call flow as shown in Figure 8. Now, after the caller has recorded his message for the station, PhonePeti asks one of the following two questions to the caller:

> *Show-timing question: Gurgaon Ki Aawaz is going to change the timings of its programs. Which program would you like to listen to at what time? Record your response after the beep.*
>
> *Women-listeners Question: We want to know which ladies from your home listen to Gurgaon Ki Aawaz. Please record your response after the beep.*

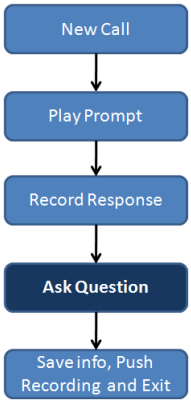

Figure 8: Call flow of PhonePeti 2 (P2)

Note that both these questions are *factual* questions [6], where the answers are based on obvious facts or awareness. Such questions are simpler to answer than conceptual questions, where answering the question requires one to evaluate several concepts.

Thus, for each call, two audio snippets were recorded: the first contained the message user wanted to record, and the second contained an answer to the question asked.

5.3 Results

P2 was deployed for one month from 7 June 2011 to 6 July 2011. During this period 438 calls were received from 258 callers. Among these, 242 had not called into *P1* earlier and were using PhonePeti for the first time. We use the 405 calls from these *new* callers for analysis to avoid bias of prior experience.

5.3.1 Improving community engagement

Even though we made several changes in the radio advertisement and prompt in *P2* to emphasize that PhonePeti is a computerized service, these efforts seem to have not had any impact on the *goodness* of calls. Out of 405 calls made in *P2*, 75.80% calls were *bad*, which is similar to 68% *bad* calls for the familiar-female prompt in *P1*. Such a high percentage of *bad* calls is a cause of serious concern as the station is losing out on valuable feedback from its listeners. A more intensive study needs to be conducted to understand the reasons behind them. Studies also need to be conducted to see if similar patterns are observed at other stations.

However, increasing the broadcast frequency in *P2* increased the absolute number of total and good calls per day. With 438 total calls and 106 good calls in one month, potential and real increment in *community engagement* due to *P2* were 45% and 11% respectively. Thus, *P2* made noticeable improvement in *community participation*; the improvement could have been significant if the percentage *bad* calls were less. This is another motivation for finding ways to reduce *bad* calls.

Providing additional instructions in the prompt to press a button to end the recording did increase the percentage of calls that were terminated by a key press. The increase from

17.62% to 25.16% was small but statistically significant (p = 0.02). Thus, explicitly asking the caller to press a button to end recordings helps, but the impact is little. Alternate ways of encouraging termination through key presses should be explored.

5.3.2 *Collecting specific information*

To collect specific information from callers, we asked them a question after they had recorded their first message. The main purpose was to explore whether an automated system could be used to solicit specific information from callers. The *Show-timing question* was asked in 208 calls out of 405, while the *Women-listeners question* was asked in the other 197 calls.

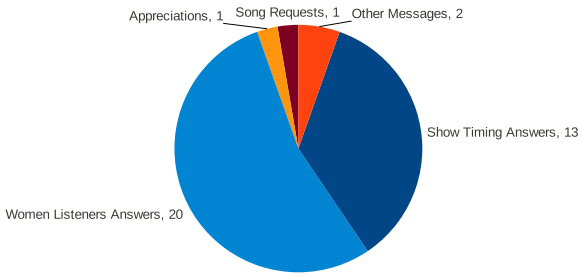

Figure 9: Distribution of second messages for calls where the first message was *good*

First message *good*: We succeeded in collecting specific information from some calls where the first message was *good*. There were 97 *good* calls, out of which 37 contained second messages. Among the rest, 29 were *empty* and another 29 were *null* indicating that the caller chose to not answer the question. The distribution of the 37 *good* second messages is shown in Figure 9. 33 of these 37 messages were answers to the specific question asked: 13 corresponded to the *show-timing question* and 20 to the *women- listeners question*. Thus, for calls with *good* first recordings, asking factual questions can solicit information from radio listeners.
First message *bad*: An interesting observation was that prompting the caller to record a second message increased the response rate. In 26 of the 307 *bad* calls, the callers went on to leave a second message. Although only 4 contained answers to the specific question asked, there were also 16 song requests, 2 program requests, 2 suggestions, 1 message for authorities, and 1 message of appreciation. This is shown in Figure 10.

It seems that the opportunity to leave a second recording made a few callers comfortable with the system, and they went on to record a message in the second try. This insight is especially relevant when we know that a large percentage of the callers call only once and give up – it may therefore be a good idea to allow the callers to record two messages rather than one.

6. SUMMARY

The results of *P2* show that PhonePeti has noticeably increased *community engagement* at GKA. In addition, there is potential to significantly increase *community engagement*

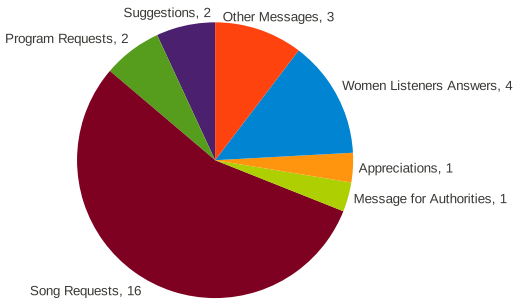

Figure 10: Distribution of recordings obtained in response to the specific questions for calls where the first message was *bad*

if the percentage of *bad* calls can be reduced. Time of day distribution of calls in *P1* show that PhonePeti is used more during off office hours, thus complementing the already existing office phone line. Content analysis of the calls in *P1* gives us a glimpse of the relationship between the station and its listeners. The results from *P2* also show that there is potential to solicit specific information from callers by asking the caller a question after he has recorded his message.

One major challenge observed in *P1* and *P2* is the high percentage of *bad* calls. We attempted to reduce caller anxiety by providing information about the system and instructions on how to use it over the radio. We also explicitly mentioned in the IVR prompt that the system is computerized. None of this information reduced the percentage of *bad* calls. Thus, ways to increase percentage of *good* calls need to be explored.

7. LIMITATIONS

A limitation of our study is that we do not have personal profile information of the callers. This makes it hard to contextualize our results within a particular demography, and thus compare our results with other related works. In addition, we have limited insight into the reasons for observed usage, particularly the reasons for such a large percentage of *bad* calls. This is mainly due to the unwillingness of one-time callers to communicate with us in telephonic interviews.

Our study also makes an important assumption: one phone number corresponds to one person. This assumption was implicit in prompt selection in *P1*, and in the analysis of the distribution of calls made per caller. We believe this assumption to be reasonable as 92% of the 411 callers who called into PhonePeti were males using mobile phones. Since mobile phones are generally perceived to be a personal device used mostly by men in the family, the phone number to caller mapping can be assumed to be one-to-one.

Our study would have been stronger if we could contrast messages and call frequency on PhonePeti with calls made to the station staff on the office phone. However, the station did not log several details of calls they received in person, making it hard for us to process the data. We are now building a simple web-interface for them to record details of incoming calls, so the we can use it to better contrast IVR systems with in person conversations.

8. DISCUSSION

In spite of its limitations, PhonePeti has demonstrated the potential for radio-telephony integration. We feel we have only scratched the surface though, and that telephony and the high penetration of mobile phones in developing regions provides tremendous opportunity for innovation in community radio. For example, an immediate extension of PhonePeti is to conduct listenership surveys. A traditional listenership survey is an expensive exercise requiring in-person interviews with a sample of the target population, or even tedious book keeping for the sample. Due to the costs involved, CR stations typically collect information from their listeners informally when they call the station to request for songs or give suggestions. PhonePeti can instead be used to conduct semi-structured surveys, either through an enhanced *P2* service, or through automated dial-outs asking short specific questions. While Lerer et al. [10] provide some insights in conducting such surveys, CR stations can utilize the radio broadcast to train listeners in answering questions in the manner expected by the system. Collecting such information can be vital for CR stations to meet their listeners' needs, and potentially even to obtain revenues by soliciting responses to advertisements and making donation requests.

9. CONCLUSION

In this study, we explored the use of an IVR system for telephony integration in the community radio context. We deployed PhonePeti, a simple answering machine system, at a CR station in India for a period of five months over two iterations. Through analysis of call statistics and call content of 758 calls made by 411 callers, we showed that PhonePeti has improved community engagement for the station. Our analysis also showed that callers learn to use PhonePeti aided by instructions aired on radio, or through trial-and-error. Finally, we were also able to solicit specific feedback from listeners, showing that IVR systems can be used to gather useful data from callers.

Several questions and concerns also emerged as a result of our study. The challenges of reducing the percentage of *bad* calls, and encouraging callers to call again, remain open. It is also unknown whether the same patterns will be observed in other CR stations. The best combination of termination criteria for recordings on IVR systems, where users do not have any past experience of using an IVR, needs further exploration. Clearly, even a technology enhancement as simple as an IVR system can require careful thought and design for its successful functioning.

Acknowledgment

The authors would like to thank Ed Cutrell, Microsoft Research India for his valuable comments on the initial design of PhonePeti. The authors also acknowledge the significant contribution of Gurgaon Ki Aawaz staff in conducting this research.

10. REFERENCES

[1] AFRRI: African Farm Radio Research Initiative. http://www.farmradio.org/english/partners/afrri/info.asp.
[2] Asterisk. http://www.asterisk.org.
[3] FreedomFone. http://www.freedomfone.org.
[4] Loudblog. http://www.loudblog.com.
[5] Sheetal Agarwal, Arun Kumar, Amit Nanavati, and Nitendra Rajput. User-Generated Content Creation and Dissemination in Rural Areas. In *Journal of Information Technologies and International Development*, 2010.
[6] H. Lynn Erickson. *Concept-based curriculum and instruction for the thinking classroom*. Thousand Oaks, CA. Corwin Press, 2007.
[7] Rochelle E. Evans and Philip Kortum. The impact of voice characteristics on user response in an interactive voice response system. *Interact. Comput.*, 22:606–614, November 2010.
[8] A. S. Grover, M. Plauche, E. Barnard, and C. Kuun. HIV health information access using spoken dialogue systems: Touchtone vs. speech. In *ICTD*, 2009.
[9] Z. Koradia, A. Premi, C. Balachandran, and A. Seth. Using ICTs to Meet the Operational Needs of Community Radio Stations in India. In *DEV 2010*, 2010.
[10] Adam Lerer, Molly Ward, and Saman Amarasinghe. Evaluation of IVR Data Collection UIs for Untrained Rural Users. In *DEV 2010*, 2010.
[11] I. Medhi, S. N. N. Gautama, and K. Toyama. A comparison of mobile money-transfer UIs for non-literate and semi-literate users. In *CHI*, 2009.
[12] N. Patel, D. Chittamuru, A. Jain, P. Dave, and T. S. Parikh. Avaaj Otalo - A Field Study of an Interactive Voice Forum for Small Farmers in Rural India. In *CHI 2010*, 2010.
[13] J. Sherwani, N. Ali, S. Mirza, A. Fatma, Y. Memon, M. Karim, R. Tongia, and R. Rosenfeld. HealthLine: Speech-based access to health information by low-literate users. In *ICTD*, 2007.
[14] S. R. Sterling, J. O'Brien, and J. K. Bennett. Advancement Through Interactive Radio. In *Information Systems Frontiers*, April 2009.
[15] Melissa Ulbricht. Press One for Freedom Fone, Press Two For Farm Radio: How Stations use Integrated Voice Response. http://www.mobileactive.org/case-studies/freedom-fone-field, 2010.