

Baang: A Viral Speech-based Social Platform for Under-Connected Populations

Agha Ali Raza^{1*}, Bilal Saleem^{1*}, Shan Randhawa¹, Zain Tariq¹,
Awais Athar², Umar Saif¹, Roni Rosenfield³

¹ITU, Lahore, Pakistan ²EMBL-EBI, Cambridge, UK ³CMU, Pittsburgh, USA
{ agha.ali.raza, bilal.saleem, shan.randhawa, zain.tariq }@itu.edu.pk
awais@ebi.ac.uk umar@itu.edu.pk roni@cs.cmu.edu

ABSTRACT

Speech is more natural than text for a large part of the world including hard-to-reach populations (low-literate, poor, tech-novice, visually-impaired, marginalized) and oral cultures. Voice-based services over simple mobile phones are effective means to provide *orality-driven* social connectivity to such populations. We present *Baang*, a versatile and inclusive voice-based social platform that allows audio content creation and sharing among its open community of users. Within 8 months, Baang spread virally to 10,721 users (69% of them blind) who participated in 269,468 calls and shared their thoughts via 44,178 audio-posts, 343,542 votes, 124,389 audio-comments and 94,864 shares. We show that the ability to vote, comment and share leads to viral spread, deeper engagement, longer retention and emergence of true dialog among participants. Beyond connectivity, Baang provides its users with a voice and a social identity as well as means to share information and get community support.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous; H.5.1.Multimedia Information Systems: Audio input/output.

Author Keywords

HCI4D; ICT4D; IVR; Pakistan; mobile phone; low-literate; telephone; speech-based social networks; blind users.

INTRODUCTION

Speech is a more natural communication medium than text for a large fraction of the world's population. Dubbed *oral* [22], such populations include people who *cannot* use text and people for whom use of speech is *more natural and preferable*. Globally, this includes 285 million visually impaired people, about 90% of whom live in low-income settings [1] and 781 million non-literate adults (13.7% of the world), 75% of whom reside in South Asia, West Asia and sub-Saharan Africa. Two-thirds of all non-literate adults are women [2, 10]. These

*These two authors contributed equally to this work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI 2018, April 21–26, 2018, Montreal, QC, Canada

© 2018 ACM. ISBN 978-1-4503-5620-6/18/04...\$15.00

DOI: <https://doi.org/10.1145/3173574.3174217>

estimates do not cover low-literates, whose numbers are far greater. Hence, non-literacy and visual impairment associate with poverty, tech-naivety and being marginalized (women, minorities, and remote). Speech is also a natural choice for cultures based on languages that do not have written forms (an estimated 46% of all languages [5]).

Most modern means of information access and communication primarily rely on textual interfaces that cannot be used by oral and tech-novice people. These people thus live mostly without modern information access, social connectivity, a voice, and a digital social identity. With a global mobile phone penetration of 73% (70% in developing countries) [8], one viable way to reach such populations is via speech-interfaces over simple phones (aka *Interactive Voice Response* (IVR) services). Even with the drop in smart phone prices, for now no other modality compares with voice over simple phones in terms of interactivity, reach, simplicity, availability and inclusiveness. IVR services do not rely on much beyond the capability of users to make and receive voice calls. As a result, users who cannot use SMS due to literacy hurdles and Internet and smart phones due to lack of availability or skill, are still able to interact with IVR services in their local languages.

However, IVR services come with challenges of their own. Such as, how to: (1) spread these services to tech-novice and poorly connected masses, (2) engage users in potentially involved voice-based interactions while keeping them oriented, motivated and grounded, (3) keep services attractive for users long term so that they keep returning, and (4) make services appealing and interactive so that users may want to contribute content. These traits of *training*, *spread*, *engagement* and *retention* are necessary for effective transmission of knowledge but are very hard to achieve simultaneously.

In this paper, we present **Baang**, an IVR-based inclusive social platform with mechanisms to achieve greater spread and uptake as well as deeper and long-term engagement built into the service. Once advertised, it keeps attracting new users while old users also keep returning to post new material and access, enjoy and assess posted content. Due to its viral and organic nature, Baang automatically overcomes user recruitment and content diversity challenges. Within 8 months, over three deployments, Baang virally reached 10,721 users who shared their thoughts via 44,178 audio-posts, 343,542 votes, 124,389 audio-comments and 94,864 shares. User surveys found that 69% of Baang's users are blind.

RELATED WORK

IVR services are being successfully utilized in several domains, including health [28], disaster response and recovery [26], entertainment [24] and social connectivity [29]. Table 1 shows a feature-by-feature comparison of Baang with existing voice-based, community forums that are oriented towards under-connected populations.

Avaaj Otalo [23], an interactive voice forum was deployed in Gujrat, India, and focused on exchange of agriculture advice among farmers. *Sangeet Swara*, a community-moderated voice forum, gathered songs, jokes and poems of up to 60 seconds (with a buffer of 10 seconds for introduction) [29]. They reported 5,376 voice posts by 1,521 callers. [25, 31] explored the use of IVR services for entertainment, job search and employment exchange. [16] presented an IVR service where community members record their songs and vote to select the best songs in a competition called *Gurgaon Idol*. [11] describe the deployment of a voice-based service called *VioKiosk* for villagers in the Juvvala Palam village in South India which links farmers, agriculture experts, and villagers, and facilitates information exchange about agriculture, health, and micro-business services. Other IVR services include citizen journalism [21], e-government [27], and marketplaces [33]. In rural India, Mobile Vaani [20] is a generalized media platform which provides many such services that run under the same infrastructure but are localized at district level. [12] used IVR to *push* informational messages to farmers and the delivery was customized through user feedback. [21] used an IVR forum to facilitate dialogue among stakeholders in a local community in India. A similar forum by [15] enabled callers in Somaliland to listen and respond to official recordings from the Presidential Ministry and other callers. [32] disseminated public health information about Ebola via peer-to-peer sharing over simple mobile phones in West Africa. [14] presented a 12-month study of a health tracking app used by community health workers that features web-based and voice interfaces in order to be inclusive towards low-literate members.

Our work, Baang, is a voice-based forum for low-literate users that *encourages diverse content, covering unconstrained genres* and also allows users to *explicitly share and comment* on the posted content. We use speech interfaces over simple mobile phones to create an inclusive, versatile and flexible platform of social connectivity in order to provide hard-to-reach populations with a voice, a digital social identity, social connectivity and community support. Compared to prior systems, Baang incorporates more features which are standard on web-based social networks like commenting on posts, sharing posts, and reporting abuse. These features have not been extensively tried or evaluated in purely voice-based interfaces (with a few exceptions e.g. [20]), in part due to a hesitation that their complexity would outweigh their benefits. As the main HCI contribution of our work, we show that the use of these features associates with high engagement and retention of users. We also document interesting social dynamics that emerge on the platform, including findings from surveys that the majority of our users are visually impaired. Taken together, these results suggest that a full-featured social networking tool

is viable and compelling on a purely spoken platform. The contributions of this work include:

- The versatile, flexible and open-ended nature of the platform that encourages users to share diverse content types and discuss a wide range of social matters rather than focusing on specific genres (see *Analysis of Posted Content*).
- An analysis of association between social media features and user engagement and retention. We show that voice-comments create dialog among participants that corresponds with deeper engagement and higher user retention (see *User Engagement with Various Interface Features, Analysis of Posted Content*).
- Uptake among blind users and females (*User Surveys*).
- The use of sharing of audio posts to achieve viral spread of the service (see *Service Uptake via Various Channels*).
- Content browsing features that allow creation of discussion threads using audio posts and make sure that popularity is not driven solely by majority votes – majority-driven content sorting poses the risk of strangling the voice of minorities (see *Use of Browsing Options*).

RESEARCH QUESTIONS

Our overarching research goal is to come up with ways to provide social connectivity to hard-to-reach populations. The success of this relies on the following specific questions:

- How can we create a simple, versatile and inclusive community forum to engage low-literate, tech-novice and marginalized (blind, women, minorities) communities and to encourage dialog among them?
- For a voice-based social platform, can we identify user interface features that associate with an increase in engagement, interactivity, spread and user retention?
- How can we encourage diversity of genres of user-contributed content?
- How can we strike a balance between the conflicting requirements of maintaining majority-driven high quality of user-contributed content yet making sure not to suppress the voice and opinions of minorities?

INTERFACE DESIGN

Baang is an IVR-based social platform for sharing short audio messages. It is accessible via phone calls and is inclusive to all types of mobile and landline phones. Baang is an Urdu word meaning *a shout* or *the call of a rooster*. While Baang can be thought of as voice-based *Reddit* (reddit.com), there are several major differences such as mechanisms of content scoring and sorting as well as modality, access and inclusion.

User Interface

A typical interaction starts when a user places a *missed call*¹ to Baang. This mechanism allows us to subsidize call airtime cost

¹A number is dialed and immediately hung up as soon as it rings; it is a familiar way of requesting a call-back in developing countries.

	Baang	Sangeet Swara [29, 30]	CGNet Swara [21]	Avaaj Otalo [23]	Mobile Vaani [7, 20]	Gurgaon Idol [16]	VoiKiosk [11]	Ila Dhageyso [15]	Polly [25]
Phone type	Simple	Simple	Simple	Simple	Simple	Simple	Simple	Simple	Simple
Medium	Voice	Voice	Voice	Voice	Voice	Voice+SMS	Voice	Voice	Voice
Mode of transmission	Broadcast	Broadcast	Broadcast	Broadcast	Broadcast	Broadcast	Broadcast	Broadcast	Peer-to-peer
Subsidized airtime	Yes	Yes	Yes	Yes	Yes	No	N/A	Yes	Yes
Subsidy mechanism	Missed call	Toll-free	Missed call	Toll-free	Missed call	N/A	N/A	Toll-free	Missed call
Vote up	Yes	Yes	No	No	Yes	Yes	No	No	No
Vote down	Yes	Yes	No	No	No	Yes	No	No	No
Report abuse	Yes	No	No	No	No	No	No	No	No
Audio comments	Yes	No	No	Yes	Yes	No	No	Yes	No
Sharing (via ID)	Yes	Yes	No	No	No	No	No	No	No
Sharing (forwards)	Yes	No	No	No	Yes	No	No	No	Yes
Length of deployment	246 days	11 weeks	2009 – now	7 months	2012 – now	33 days	4 months	5 months	1 year
Calls	269,468	25,000	137,000	6,975	10,000/day	306	20,499	N/A	636,000
Users	10,721	1,500	715+	45	1,500,000+	252	976	N/A	165,000
Posts	44,178	5,000	13,595	896	300/day	31	2,532	4,300	387,301
Female Users	10%	6%	21%	0%	N/A	Low	N/A	15%	11%
Visually impaired users	69%	26%+	N/A	N/A	N/A	N/A	N/A	N/A	Less than 1%
Content genre	Open	Fixed	Fixed	Fixed	Open	Fixed	Fixed	Fixed	Open

Table 1. Comparison with other services (N/A means that the particular statistic was not reported)

for the users, hence encouraging longer and more engaging interactions. As soon as a (missed) call is placed, a ring-back-tune informs the user that Baang will call back soon. Baang calls back within a few minutes. If users do not answer the return call, Baang retries after 2 minutes then again after 5, 10 and 20 minutes before giving up on the request.

As soon as users answer the return call (Figure 1), Baang plays culturally appropriate greetings followed by a disclaimer/warning that users should record content politely and responsibly and that recordings containing foul language would be deleted and users repeatedly posting such content would be blocked. Users are also informed that any content they record would be available to public and will also be used for research purposes. Next, Baang is introduced as:

Now you can share your baangs, such as <content suggestion> with the rest of the country.

In order to build the very image of the service around content diversity, the suggestions within angled brackets are randomly selected from five options: (1) news, interesting anecdotes or a nice thought, (2) recitation of Quran or naat (religious poetry), (3) a song, poem, or ghazal (a famous genre of poetry), (4) a folk song or folk tale, (5) problems of your locality.

Next users are given three options: (1) record a new voice post (baang), (2) listen to baangs recorded by others, or (3) to check the status (votes, comments) of their own previously recorded baangs. Users can select these options via key presses. In case of an invalid input or no key press, the last voice-prompt is repeated and users are allowed to retry.

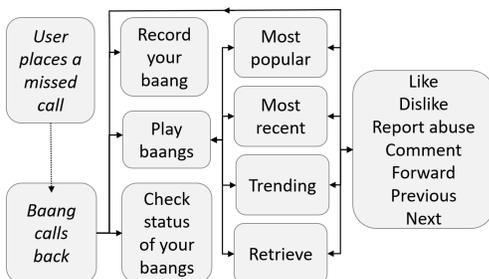


Figure 1. High-level call flow of Baang

Users who choose to **record** a new baang are given up to one minute of recording time. Users are reminded to record content politely and responsibly and make sure not to disclose any personal information like addresses, names, and phone numbers. The recorded baangs are played back to the users for confirmation before getting posted. After recording their first baang, users are also asked to record their names that are announced as part of the baangs that they share with their friends. Users are also provided a 4-5 digit unique identifier for their baang that could be used to retrieve it and are encouraged to advertise it to their friends to gather votes and comments.

Users, who choose to **listen** to baangs recorded by others can browse by recency (newest to oldest), popularity (most to least up-voted baangs over last seven days), and trending (most to least up voted baangs over last 24 hours). Users can also retrieve a particular baang by entering its ID. The score (# up-votes) of each baang is played before it. After each baang, users have the options to: (1) navigate to the next baang, (2) previous baang, (3) like, (4) dislike, (5) report abuse, (6) post a comment (that also allows them to listen to the existing thread of comments (newest to oldest)), or (7) share the baang with friends (by entering their phone numbers). A user can like, dislike, or report-abuse for a particular baang only once. The dislike and report-abuse scores are not used for sorting content. We use these scores to flag potentially inappropriate content to be scanned by human moderators (as explained in section *Analysis of Posted Content*).

Users who choose to **check** status of their own previously recorded baangs have the option to listen to the number of likes, number of dislikes, or the audio comments received by their baangs. They can also delete comments posted on their baangs.

Interface Design Methodology

This section discusses the reasons behind some of our interface design choices. We have designed the interface to be simple to comprehend and easy to navigate. All menus and recordings can be skipped using # to accommodate advanced users. Prompts are recorded by a male voice artist in a friendly voice. Simpler options are kept towards the start of menus. To enable **content browsing**, community forums sort posts by balancing likes and dislikes so that the most sought-after

material makes it to the top (e.g. [29]) and new material gets adequately exposed. We give our users the choice to listen to recent, popular or trending posts. The popular and trending posts are based on up-votes only. Following are the reasons:

1. A majority-driven popularity risks suppressing the voice of minorities. Posts pertaining to problems, feelings and views of minorities especially ones that conflict with the will and interests of the majority do not get up-voted and end up getting buried in the play-list. It is necessary to provide an alternate yet fair way of accessing such content.
2. Sorting schemes that weigh-in likes and dislikes may be "fair" but are not easy to explain to low-literate users. Simple sorting based on up-votes is easy to explain though it misses out on down votes. Our focus has been to make Baang simple and accessible to tech-shy users, therefore we opted for up-votes-based sorting. We also play the score of each baang before it to ensure transparency. We never received any complaints about content sorting from users in their otherwise very vocal feedback.
3. Competitors manipulate votes. One user overtakes another by either promoting his own baangs or by asking his friends to down-vote content posted by his competitors. We found evidence of this behavior as discussed in *Content Analysis*.
4. Users respond to content posted by others by commenting as well as posting baangs. A chronological listing option is required to follow such threads.
5. The sliding windows of seven days (arbitrarily chosen) and 24 hours in popularity and trending lists prevents monotony and keeps the top content rolling.

The features of **sharing** and **retrieval** of baangs by ID were added to obtain organic spread of the service. As users spread word about their content, they effectively spread the service virally. We added **voice comments** to allow users to leave rich feedback on posts. This allows users to interact in a more personal manner and to have discussions and dialog. To keep the interaction simple, we allow users to navigate comment threads and post their own comments via a brief sub-menu. The first deployment of Baang just allowed users to like and dislike posts. We later added the **report abuse** feature to allow users to flag inappropriate content when we found cases of foul language. **Genre suggestions** were added to promote content diversity. We wanted Baang to become a virtual community conducive to diverse and rich interactions rather than limiting it to a particular genre or domain.

Seeding and Launch

Baang was deployed in Pakistan in 2015. With a high teledensity of 70% [4], and a low rural adult literacy rate of 49% [3], there is a clear case for using inclusive means of information dissemination and social networking in Pakistan. Baang was deployed thrice, each time targeting either a specific seeding and advertisement mechanism or usage quotas on subsidized usage. Figure 2 shows the number of calls, users, and new users per day for all three launches (discussed in detail in the next section). Baang was initially launched on Nov 02, 2015 for 41 days. This version did not have the sharing via

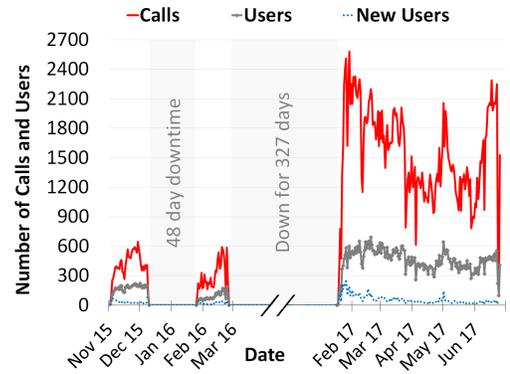


Figure 2. Calls, distinct users, and daily new users

forwarding capability and voice comments were added a week into the deployment. The key features of this launch included: (1) continuous advertisement over Polly, a popular IVR service in Pakistan [25] (described in section *Spread and User Retention*), (2) line capacity of 30 simultaneous calls, and (3) usage quota of 2 free calls per user per day. Users could call a paid line as much as they liked. Usage quota was introduced to manage call traffic so that new users get return calls from Baang reasonably quickly.

After the first launch, Baang remained down for 48 days as we made some major upgrades to the telephony backend. We ported Baang from a commercial telephony platform, Tropo [9], to an open source platform, FreeSWITCH [6]. A major change in the second deployment was that we stopped all seeding and advertisement. We did not even advertise on Jan 26, 2016 when we switched Baang back on. As we kept the phone number same between the two deployments, users discovered on their own that Baang is back, and started calling. We removed the paid line as it had only received 454 calls from 204 users in the first launch. The second deployment lasted 31 days.

Baang remained down for 11 months as we acquired funds to support a longer deployment. It was made live again on the same phone number, on Jan 04, 2017, without any announcements or advertisements. Surprisingly, we started getting phone calls immediately. This deployment lasted 29 weeks and featured an increased call capacity of 90 concurrent calls that also allowed us to remove the usage quotas. We announced the approaching end of deployment to our users and shut Baang down on Jun 26, 2017.

RESULTS AND DISCUSSION

Call Volume and User Engagement

We find very high and active user engagement. The three deployments spanned a cumulative of 246 days (approx. 8 months). Over this period Baang received 269,468 phone calls from 10,721 users and a total of 44,178 posts that were contributed by 2,990 users (for comparison with other platforms, see Table 1). Of these, 36,986 posts were well formed (non-zero recording duration) that constituted 26,253 minutes (438 hours) of audio data. The posts were listened 2,817,322 times (an average of 11,890 times per day, 76 times per post); voted up 248,227 times by 4,486 users; voted down 54,790

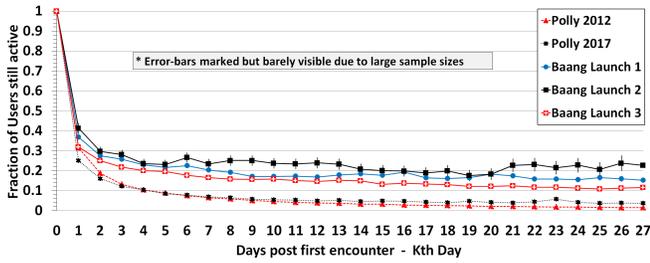


Figure 3. User Retention

times by 2,369 users; and reported for abuse 40,525 times by 1,915 users. A subset of baangs was shared (forwarded) 94,864 times by 2,350 users with 7,574 friends. After Baang was shutdown, we turned on the caller-paid line for 14 days. We received 1,088 calls from 463 users while 226 users called more than twice. The average call volume went down from 779 calls per day for the subsidized lines to 78 calls per day for the caller-paid line. This reduction is consistent with the findings of Vashistha et al. [29] and Raza et al. [25].

Baangs received 124,389 voice comments from 2,490 users. The comments constituted a 38,755.6 minutes (646 hours) of audio data. Number of comments are 2.8 times more than the number of baangs and 1.47 times greater in terms of average recording time. The comments to likes ratio of Baang is 1:2.

Call traffic on Baang consumed 1,781,271 minutes of airtime costing PKR 2,811,151 (USD 26,689). On average, this means that every user spent 2.8 hours on Baang with each call to baang being 6.6 minutes long. Each baang costed PKR 63.63 (60 cents). In terms of *stickiness*, defined as the ratio between *Daily Active Users*² (DAU) and *Monthly Active Users* (MAU), Baang steadily climbed from an initial 13% to 18%. Average sessions per DAU were around 3.25 calls.

Use of Browsing Options

As discussed, the 36,986 baangs were played 2,817,322 times. Of these, 60% of all play instances were through the recency listing (by 3,276 users); 9% via popularity (by 2,167 users); 26% via trending (by 3,319 users); while 5% via Baang ID entry (by 6,360 users). These numbers indicate a clear user preference towards recency listing, followed by trending. In addition to reasons discussed earlier, these were also the only two lists where return users are most likely to find fresh content. Also, the highest single fraction (15%) of all baangs comprise discussions among users that could only be browsed in order using recency listing.

Spread and User Retention

Figure 3 shows the user retention of the three launches of Baang compared against the user retention of a popular IVR-based, viral entertainment service, *Polly* [25], that was deployed in Pakistan in 2012 (for a year) and again in 2017³. We perform this comparison to provide more context to our findings. Polly virally disseminates development-related information to low-literate populations using entertainment as a hook. It engages users by allowing them to record and morph

²An active user on any date is defined as one who interacts with Baang on that day by placing a call or receiving a forwarded message.

³This data is being reproduced with permission from Raza et al. [25]

their voice with funny voice modifications and to forward these to friends. Once engaged, users are exposed to various useful services like job search [25] and health information [32]. Seeded with 32 initial users in Pakistan, Polly organically spread to 165,000 users over a year, who took part in 636,000 calls and 33,000 also started using job search.

In Figure 3, each plot shows the fraction of users who continue using the service k days after their initial interaction (where $k=1$ to 27 in Figure 3). Use of service is defined as at least one call on a particular day (k days after user's first interaction). The denominator only counts the subset of users who had a chance of using the service on their k^{th} day. For example, a user who starts using baang two days (48 hours) before the end of a deployment only had a chance to use Baang for two days. Confidence Intervals are marked on each plot, which show that all these differences are statistically significant. Because of the large sample sizes, the confidence intervals are barely noticeable. For readers more familiar with *churn* or *attrition* plots, $retention(k)$ is defined as $1 - churn(k)$.

In Polly, 24% (Polly'12) to 31% (Polly'17) of users return on their second day and 16% to 19% on their third day. The ratios drop below 7% after a week and 4% after two weeks. It is also notable that a small yet significant fraction of users (0.5% to 1%) keep returning to Polly for several months (not shown). In the quota-restricted deployments of Baang (1 and 2), around 40% of users return on their second day while 28% return on their third day of use. In the unrestricted version (Baang 3) 32% return on their second day, while 25% return two days after their initial interaction. This shows that the **retention is significantly high as compared to Polly** (notice the non-overlapping confidence intervals). However, the differences become more apparent after a week when more than 17% to 20%, and after two weeks when 15% to 19% of users return to Baang. Between 10% to 20% of Baang's users (compared to 1% to 3% in case of Polly) keep returning after four weeks.

There is a clear difference between the first two, and the third launch of Baang both in terms of uptake and spread (Figure 2) as well as user retention (Figure 3). Number of calls and users per day climbed rapidly right from the beginning of the third deployment of Baang to nearly four times that of first two launches. This could be attributed to the removal of usage quota of 2 subsidized calls per day. However, this increase in traffic coincided with a drop in user retention (Figure 3). Our hypothesized explanation is that only more committed users end up calling back the quota restricted model of Baang (primed with only 2 free interactions on their first day). With no usage quotas, many users just call to test Baang without committing to its usage and do not end up becoming long-term users. Such users churn at a higher rate. Another explanation could be that the quota-restriction makes Baang more attractive and tantalizing to users as opposed to its free version.

User Engagement with Various Interface Features

Figure 4 shows the overlap among sets of users who engage with various interface features of Baang, using an *upset graph* [18]. The graph shows sizes of sets of users who engage with various social media features and non-zero overlaps (intersection) among these sets. We find that 1,609 users (26% of all

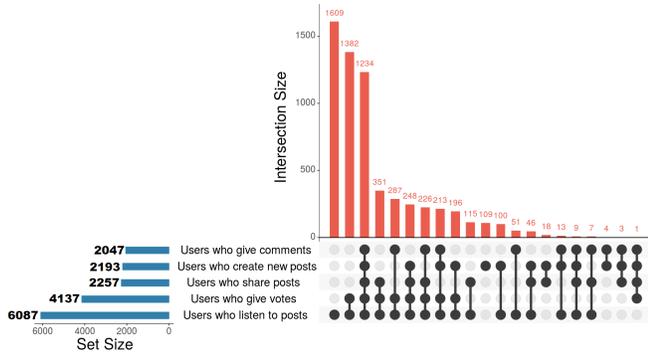


Figure 4. Overlap among users who engage with various features

users) just listen to posts without any other actions, followed by 1,382 (22%) users who listen to posts and vote on them, and interestingly, a comparable **20% of all users engage with all of Baang’s features**. The rest of this section explores the behavioral traits of these users who engage with various features of Baang in terms of their uptake, engagement and retention.

Engagement

In this section we explore the interplay of various social network features and users’ engagement and retention. Our hypothesis is that the users who engage with more sophisticated and proactive features where they are required to provide actual input (voting, commenting, posting, and forwarding as opposed to passive listening), are likely to engage in lengthier and more frequent interactions (engagement) and to keep returning to the service for longer periods of time (retention). Users who post baangs and comments are more likely to call back (to find out the number of responses, votes and comments received by their posted content) as compared to users who never post audio. Such users are also likely to spend more time in calls as posting content is a time consuming activity.

To find evidence for the above-mentioned, we analyze how the distribution of number of calls, call lengths and number of days for which users remain active associate with engagement with various features of Baang (comment, vote, post, and forward). As all of these three engagement metrics, follow a Zipfian distribution⁴ across users, we find it more expressive

⁴The call frequency of a user is inversely proportional his rank in the frequency table. This is also true for call lengths and active days.

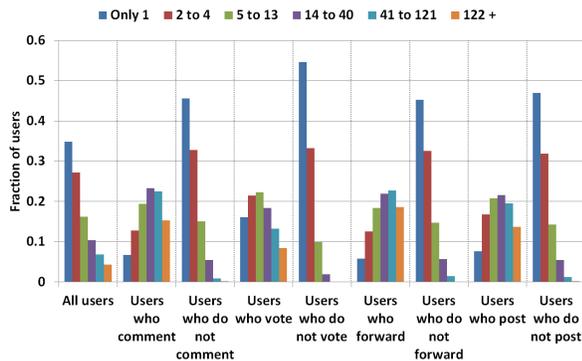


Figure 5. Number of calls (by user groups)

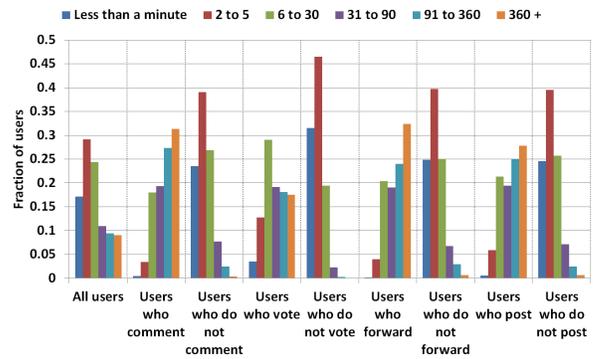


Figure 6. Total duration of calls (by user groups)

to visually compare the distributions side-by-side. To better visualize the distributions, we use bucketing with exponentially increasing bin sizes. The actual bin sizes were chosen after visually experimenting with various exponential functions.

Figure 5 shows the distribution of calls across users (normalized by number of users in each group). The leftmost distribution is for all users. We find that 35% of users just perform one call to Baang and hence do not engage at all. This is followed by 27% of short-term users who perform 2 – 4 calls, 16% of medium-term users who perform 5 – 13 calls, 10% of long-term users who place 14 – 40 calls and 12% of "Baang Junkies" who place more than 40 calls.

Next, we compare across user groups. These groups are pairwise mutually exclusive and hence pairwise comparable e.g. users who comment vs. users who do not comment. We find that in each pair, the distribution of calls (Figure 5), call lengths (Figure 6) and number of active days (Figure 7) over users is skewed more towards the right (more users with a high number of calls) for users who *engage* with various interface features as compared to users who *do not engage*. We find the highest right skew (engagement with Baang) for user groups who engage with comments and forwarding features and the highest left skew (lack of engagement with Baang) for users who do not even engage with voting. Figures 7 and 8 show that majority of the users who do not engage with these features stop using the service within a week. In fact, **63% of users who do not even engage with voting stop using Baang just after one day of use**. We performed Mann–Whitney *U* test [19] to compare user groups (e.g. commenters vs. non-commenters)

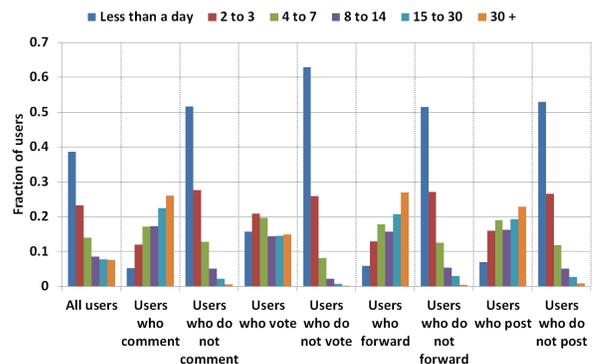


Figure 7. Number of active days (by user groups)

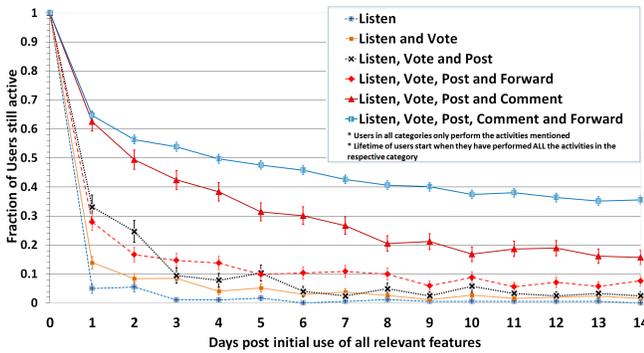


Figure 8. Comparison of retention across users after they engage with various interface features

over each feature (number of calls, length of calls, number of active days) and found the differences to be statistically significant (p -value < 0.001). User **engagement** (frequency and length of calls) and **retention** (number of active days) is highly associated with use of more sophisticated interface features like comments, sharing of posts (forwarding), creation of posts and voting. However, these results do not establish a causal relationship between features and engagement but only points at high degree of association.

Retention

We next compare some of the interesting user groups from Figure 4 for retention. Our hypothesis is that users who engage with more of Baang’s interface features are also likely to be the ones who end up calling back for several days. Figure 8 shows the comparison of retention among six mutually exclusive groups of users. We have conservatively measured retention across users who engage with various features of Baang *after* they have performed *all* of the actions stated for that group at least once (making it their day 0). For example, for the group "Listen, Vote and Post", the day a user finishes all these three tasks at least once, we consider that to be his day 0. Next, we find the retention across these groups as a function of days just like we did for Figure 3. Again, confidence intervals are marked on each plot that show that these differences are statistically significant. Because of the large sample sizes, some confidence intervals are barely noticeable.

We find that indeed the group that performs worst in terms of retention is passive users i.e. users who listen to baangs but do nothing else. The group that performs best is the one that engages with all of Baang’s features. Interestingly, more than 36% of the users in this group keep calling back for up to two weeks. The only difference between this group and the next lower plot is that users in the lower group do not share posts with friends. However, we also clearly see the significance of comments as the next lower group is the one where users listen, post, vote and forward but do not comment. The three groups that are lowest in terms of user retention are the ones where users do not engage with comments and forwarding. We also find it interesting that **the act of posting Baangs alone does not correspond with very high user retention.**

Predictive Termination of Use

Next we explore if engagement with interface features could be used as predictors for continuation (or imminent discontin-

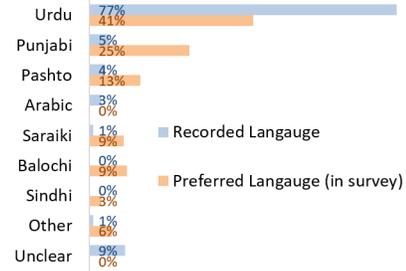


Figure 9. Language distribution in Baang & preferences

uation) of use. We have based this analysis on the simplest of questions: "How likely is a call to be the last call from a particular user if he performs an activity (like commenting, voting, forwarding or posting) in that call vs. if he does not?". Thus this question explores the short-term impact of engagement with an interface feature. We use Maximum Likelihood Estimation with first-order Markov chains. We find comments to be the best predictors of discontinuation-of-use. The likelihood that a user will never call again if he comments in the current call is 0.0140 (Confidence Interval [0.0134 – 0.0146]), vs. 0.0450 (CI [0.0449 – 0.0459]) if he/she does not. Next most informative feature is forwarding. Interestingly **creation of new posts is a worse predictor of continuation of use** (0.0358 if posted vs 0.0481 if not) as compared to voting (0.026 if voted vs 0.0478 if not).

Service Uptake via Various Channels

Baang has several spread mechanisms. We find that 60% of all users had their first interaction with Baang via calls that convey forwarded baangs to their friends, 12% had their first interaction with Baang via Polly (our advertisement allowed users to press a key on Polly to transfer to Baang). Users could also inform their friends about Baang in offline conversations (Word-of-Mouth (WoM)) and also advertise IDs of their recorded baangs and ask them to promote these posts. 28% of all new users came in through WoM out of which 8% retrieved a baang by entering its ID in their very first call to Baang. Therefore, **the most successful mechanism of spread in Baang responsible for bringing 60% of all users was forwarded posts**, followed by WoM, advertisement over Polly and retrieval of baangs by ID. We also find that the users who came to Baang via forwarded posts and later started calling Baang on their own were the ones who performed best in terms of engagement and retention (analysis not shown).

ANALYSIS OF POSTED CONTENT

Content Annotation

All posted content is regularly monitored by two paid moderators. Each has more than fourteen years of education and one of them is a trained linguist. They annotate all recordings for gender (male/female/unclear), profanity (inappropriate/appropriate/unclear), language and genre. To prevent misuse of the service, they immediately listen to all posts that are reported for abuse by users. Inappropriate posts are immediately removed and a warning message is played to the users responsible for contributing them (in their next call to Baang). Repeat offenders are blocked from Baang.

Genre	f	%
General		
- Discussions, arguments and grievances	1,962	
- Greetings (hello/hi, special event/birthday)	786	
- Advertisement (personal and business ads)	604	
- Self Introduction	497	
- Sayings and quotes	450	
- Information (Science, general)	329	
- Nice thoughts	147	
- Posting ethics	92	
- News	85	
- Sports	72	
- Feedback about Baang	56	
- Jokes	45	
- Voice Selfies and Feelings	30	
- Local area problems	4	
<i>Subtotal</i>	5,159	39
Religious Content		
- Poetry and hymns	1,418	
- Information, quotes, prayers	861	
- Recitation of Quran, Azan	473	
<i>Subtotal</i>	2,752	20.8
Songs and Poetry		
- Songs	788	
- Poetry	669	
- Folk songs	73	
- Patriotic songs	23	
<i>Subtotal</i>	1,553	11.7
Empty		
- Silence	715	
- Noise	209	
<i>Subtotal</i>	924	7
Asking for prayers and support	707	5.3
Unclear	681	5.1
Hosted Programs	623	4.7
Inappropriate Content		
- Profanity	540	
- Hate speech	30	
<i>Subtotal</i>	570	4.3
Questions		
- Quiz questions (sometimes with airtime rewards)	66	
- Quora-like questions	60	
<i>Subtotal</i>	126	1
Prerecorded Content		
- Music	68	
- Folk music	46	
<i>Subtotal</i>	114	0.9
Other	19	0.1
Total	13,228	100

Table 2. Distribution of genres

For genre, we use hierarchical categorization (using open coding). The main genres are shown as bold in Table 2, along with their sub-genres and the frequency (f) of each sub-genre. For simplicity, the annotators are asked to initially classify each recording by the main genres before tagging the sub-genres. For classification of profanity, annotators are instructed to classify each recording as *inappropriate* if it contains at least one inappropriate word (e.g. foul language) or hate speech. To verify agreement, a random sample of 103 recordings was classified by both annotators. The inter-annotator agreement around genre was found to be $\kappa = 0.97$ as measured by Cohen's kappa, which is defined as $\kappa = \frac{P(A) - P(E)}{1 - P(E)}$, where $P(A)$ is the proportion of agreement and $P(E)$ is the proportion of agreement by chance [13]. Despite the absence of any reference lexicon of profane words, the agreement around inappropriate content was found to be $\kappa = 0.88$. Annotations of the gender and language attributes showed perfect agreement. According to Krippendorff's interpretation [17], these kappa values mean that our annotation is reliable.

Analysis

These stats are based on 13,228 randomly selected baangs. Most of the recordings (91%) contain actual content while only 5% constitute noise and silence (remaining 4% are hard to categorize). Users tend to report abusive content and enthusiastically compete for votes. **Only 4.4% Baangs contained profanity.** Baangs are mostly contributed by males (90% of baangs) with a small yet significant fraction (10%) of female-recorded posts. It is notable, that **among the 100 top liked posts, 14 were contributed by females.**

Figure 9 shows the distribution of baangs by language. Urdu, the most widely understood language in Pakistan, is used far more than any other language. Urdu is also the interface language of Baang which could also be a factor contributing to its greater use (our surveys show that the users themselves come from diverse linguistic backgrounds). The Arabic recordings comprise of recitation of Quran and other religious content. Code-mixing is a not very frequent with only 1.8% posts where users mix two or three languages.

There is a great diversity of genres as shown in Table 2. The most popular sub-genre "discussion, arguments, and grievances" spans casual gossip and discussions around current affairs, social problems, personal matters and posted content. Also included are directed comments and feedback regarding quality and content of specific posts, and heated arguments around matters of mutual significance like posting ethics. The "religious" categories includes some of the most highly up-voted baangs. 15% of all posts also contain requests for votes, where people advertise their posts and personality and appeal for votes and comments. Prerecorded content comprises recordings of songs and music being played on other devices.

As discussed in section *User Interface*, Baang also gives suggestion regarding genres. 36% of all baangs belong to the suggested categories. In 21% calls, users recorded genres suggested to them in that very call.

Interesting and Unexpected Content

Baang is put to several interesting and unexpected uses. Community support circles rapidly form and people raise their voice against disability abuse and female harassment. Users taunting or making fun of the blind are collectively criticized, reproached and counseled. Harsh comments or rude language against females face strict opposition from both male and female users. In 5.3% of baangs, people asked for emotional support, advice, sympathy and prayers. These requests were met very favorably. We found 60 instances of Quora-like questions (quora.com) where users ask for instruction about device repair, accessibility software bugs, and questions pertaining to everyday activities (e.g. questions about specific NGOs). People also ask for counseling regarding matters like love and breakups. These questions were answered through comments and follow-up posts. Among discussions a prevalent theme is posting ethics (what should and should not be recorded). Users differ around matters of posting personal details like addresses and phone numbers, liberal and conservative content, intimate and explicit material that is not suitable for children etc. Majority of users unite against foul language, hate speech, harassment, and disability abuse. People also discuss current

affairs, vent out against terrorism, call for national unity, and show support for the forces that are fighting terrorism.

A popular trend on Baang is **radio-style, sponsored, hosted programs and dramas and plays** that comprise 4.7% of all posts. A program or drama is recorded in the form of a series of baangs. These programs have catchy titles. One such program is hosted by two users who take calls and SMS requests for songs and poetry recitals, interview guests, pose quiz questions with (100 rupee) airtime prizes. It is sponsored by a skill training organization for blinds. 4.6% of all Baangs are advertisements. Users sell products, advertise their services, businesses and organizations. For instance, a special education teacher asked blind people to join his school, a person offered to teach English Language classes on Baang, a government official offered accessibility services for the blind, users even planned protests and kept posting updates to their plans. Some users also posted personal ads for friendship and marriage. Baang also gives people a chance to share their art and creativity. Poets recite their poetry, preachers deliver sermons and lectures spanning across several episodes.

Voice selfies and feelings emerged as another interesting trend. We found 30 posts in our sample where people posted their activities and feelings, usually with a lot of details. A user posted in the middle of the night that he is having biryani (a rice dish) and even recorded munching sounds and detailed description of the flavor. Another posted about feeling happy at a friend's wedding and expressed his excitement through exclamations. There were also "thank you" and apology posts. A large number of posts are specific to **blind users**, their problems and rights. Lobbying for support to get a quota in the National Assembly, uniting against misbehavior, raising voice for their rights (wages, getting accounted for in census, opportunities for jobs, education, and skill training), awareness campaigns about rights of people with special needs and their role in development, marriage bureau, audio books, dedicated forums and advertisement of services like housing and food for low-income blinds. Several posts educated people on how to address and talk to the visually impaired.

Among Baang's negative uses is foul language and hateful speech (4.3% of all posts). It significantly reduced when we started playing users' recorded names along with their baangs and comments, and started removing inappropriate posts and blocked repeat offenders.

Top Ranked Content

We annotated all baangs that had been played more than 200 times each (N = 1,981) and calculated the ratios of likes, dislikes and reports to total votes (likes+dislikes+reports) for each baang. We analyzed top 100 baangs of each category. None of the **100 top-liked baangs** were empty or contained inappropriate content. Recordings had high quality audio, sometimes with sound effects (echo, music). There were 14 female and 86 male recordings with 79 Urdu, 2 Punjabi and 19 Arabic (Quranic recitation) posts. Only one of these recordings contained requests for votes. 80 posts belonged to religious category (poetry: 33, prose: 29 and recitation: 18), followed by 16 sayings, anecdotes and quotes, while the remaining four contained greetings, and poetry.

None of the **100 top-disliked baangs** were empty, and most (N = 95) were civilized. There was higher linguistic diversity with 89 Urdu, 4 Punjabi, 3 Arabic, 2 Pashto, 1 Saraiki and one prerecorded baang. There were 47 female and 53 male recordings and 16 had requests for votes. There was more diversity of genre with discussions (27), songs (18), religious posts (15), poetry (11) and sayings (10) as the main categories. These posts received a high fraction of dislikes not because of low quality or profanity but because (1) competitors get their opponents' content down-voted, and (2) conflicting values (liberal vs conservative; with each group down-voting content of the other e.g. response towards religious content vs. singing and music). This supports our decision to exclude dislikes from content scoring.

The **100 top-reported posts** indeed represented highly inappropriate and offensive content. None of them were empty. Three of the most reported baangs contained pornographic audio clips. Of the remaining 97: 32 were posts in which users were yelling or having serious arguments (bordering on profanity), 26 contained abusive language, 4 threats, 5 sectarian hatred, 5 where users (1 male, 4 female) shared their cell numbers and asked people to call them, 2 with allegations against people and NGOs, 2 with people making fun of other users. The remaining 21 otherwise civilized posts just happened to be recorded by users who were engaged in fighting others. There were 95 Urdu and 2 Punjabi recordings. 19 were recorded by females, 78 by males while 3 contained prerecorded content.

Analysis of Audio-Comments

We analyzed 1,200 audio comments using the same coding scheme that we used for baangs. As compared to baangs, comments contain more empty recordings (12.5%), less linguistic diversity (six languages with 71% Urdu and 9% Punjabi recordings) and a higher proportion of foul language (14.4%). Female participation was a bit low in comments (6.3%) as compared to baangs (10%). 21% of comments were discussions, arguments and exchange of ideas, mostly around the baang on which they were posted while another 25% were greetings, introductions and casual hello/hi messages. 15% of all comments comprised of discussions on posting ethics. These fractions are significantly higher compared to the corresponding fractions in baangs. This shows that comments were used to create discussion threads where people exchanged ideas and opinions. There were very few recordings with songs (0.42%) and poetry (0.67%). Apparently users prefer to post such content in the form of baangs for higher visibility.

User Surveys

We conducted telephonic surveys of 500 users, randomly selected from among the 10,721 users of Baang, to collect demographic details. Out of these, 276 users consented to answering our survey questions. We asked them for their name, age, number of years of education, location (district), profession, brand and model number of mobile phone (we used this to determine simple/feature or smart phone), internet access (yes/no), monthly mobile expense, if they had any visual impairment and the language they use with their family. The surveyors also noted the gender of the participants. Open ended questions included: reasons for using Baang, whether they would still use

Baang if not free and any feedback and suggestions. Free-form responses were annotated offline using open-coding.

Survey participants were **mostly low educated, young men** with a large fraction (13%) having no formal education (see Figure 10). However, there were also a significant number (32%) with more than 10 years of education. The participants belonged to 85 districts of Pakistan across all provinces (including Baluchistan (2%), Gilgit-Baltistan (1%), FATA (0.5%) and AJK (1.5%)). Their majority was from Punjab (69.5%) and Khaiber-Pakhtunkhwa (15%). Of the 262 that answered regarding their professions, **30.5% were unemployed adults**, 19% were students, 11% just said that they were employed, followed by 8% government employees, 6.5% teachers, 5.4% farmers, 4.2% manual laborers and a handful of mostly low-educated professions like shopkeepers, drivers, carpenters, cobblers, gardeners, house wives, shepherds, tailors, cooks, electricians, factory workers, police, and community health workers. 75% (N=276) of surveyed users owned simple or feature phones, 21% had smart phones, while 4% owned both. 56% (N=262) did not have access to internet. **65% of 242 users reported a monthly mobile expense of less than PKR 500 (USD 4.7)**, and 83% reported less than PKR 1,000 (USD 9.5). We asked them if they would still use Baang if the air-time is not subsidized, to which, 46% responded yes, 34% responded no, 8% did not respond while the rest (12%) said that they will use it if partially subsidized or if they could find bulk airtime at low cost. Notice that despite these stated preferences, average call volume went down from 779 to 78 calls per day after we removed airtime subsidy. However, some users appreciated it as they believed that free access encourages non-serious and profane content. A male user recorded:

"Hello Baang team! I'm happy that you came to your senses and started charging. It is a good step. Now whoever curses, will curse at their own expense! They (people who curse) behave as if their father owns this service..."

Another male user recorded:

"... I am happy that charges have been imposed. Now no one will record curses, foul stuff, and sectarian hatred. My advice is that Baang should also require users' ID card numbers, so that people refrain from bringing bad name to the service..."

Surprisingly, 69% (N=189) of the people that we contacted were visually impaired. This was a major finding as we had not advertised Baang to the blind community. Blind users praised Baang profusely. Apparently, most of them assumed that Baang is designed specifically for visually impaired users. A blind user recorded:

"... I'm blind and inexperienced... This is my first experiment. So whoever listens to my voice should reply. I'll keep recording baangs. I do not know what to say... Keep me in your prayers..."

Another interesting finding was regarding preferred language (Figure 9, N=268): Although Urdu dominates among recorded languages, **our users belong to diverse linguistic backgrounds and prefer communicating in their regional languages.** The fact that Urdu is the interface language of Baang, may have prompted mostly Urdu recordings.

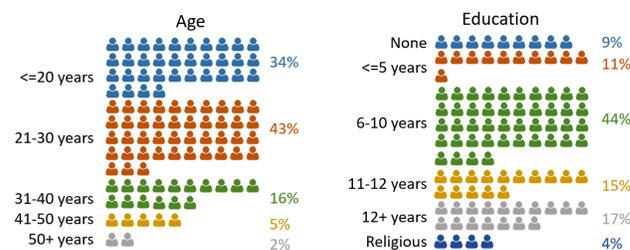


Figure 10. Distribution of age and education in survey

Most users praised Baang but several complained about profanity. We received interesting responses on how Baang improves their life. Blind users defined Baang as an alternate source of information, ads, news and entertainment to books and newspapers (low-income, blind users cannot pay others to read books to them), making friends, discussions, job search. Among their suggestions were to filter foul language, increase recording time, audio books, friend-lists for sharing baangs within specific groups, keeping the service free for blinds, sports shows, prizes for high vote achievers. Some users apologized for foul language and asked us to unblock their numbers.

CONCLUSION AND FUTURE WORK

We presented Baang, a versatile, flexible, and inclusive voice-based social platform for hard-to-reach and oral populations. Analysis of gathered data shows that Baang creates a vibrant community of users from diverse socio-economic and linguistic backgrounds including 69% blinds, 10% females and mostly low-educated, unemployed, young men from all over Pakistan. Baang's open community included people from remote areas and linguistic minorities. Social network features like content sharing and voice comments led to viral and enthusiastic uptake of the service, high user engagement and retention, and true dialog among the community. Browsing and scoring mechanisms of Baang ensure majority-driven quality assurance but not at the risk of drowning the voice of minorities. Baang provides a window into the collective values of a community as they raise their voice against disability abuse, female harassment, foul language, hatred, terrorism and unite for their rights and in support of the oppressed. We show that voice-based social platforms can provide under-connected and tech-naive individuals with a voice and social identity. Next, we plan to annotate the speech content for various interesting features like prosody, accent, sentiment. This would enable localization of linguistic resources for Pakistan. We also plan to analyze Baang's social network dynamics.

ACKNOWLEDGMENTS

We thank ITU for supporting this work, our reviewers for their insightful feedback that improved our paper a lot, our research assistants Hira Ejaz, Abdullah Sabri, Gulnar Khan, Arooj Saleem and Mukhtar Ahmad for spending lengthy hours annotating data, moderating posts and performing surveys.

REFERENCES

2014. WHO | Visual Impairment and Blindness. <http://www.who.int/mediacentre/factsheets/fs282/en/>. (2014).

2. 2015. World Bank - Literacy rates. <http://data.worldbank.org/indicator/SE.ADT.LITR.ZS>. (2015).
3. 2016. *Pakistan economic survey 2015-16*. Technical Report. Economic Adviser's Wing, Finance Division, Government of Pakistan, Islamabad.
4. 2016. *PTA annual report 2016*. Technical Report. Pakistan Telecommunication Authority.
5. 2017. Ethnologue: Languages of the World. (2017). <https://www.ethnologue.com/country/US>.
6. 2017. FreeSWITCH. <https://freeswitch.org/>. (2017).
7. 2017. Gram Vaani. <http://www.gramvaani.org/>. (2017).
8. 2017. GSMA Mobile Economy 2017. <https://www.gsma.com/mobileeconomy/>. (2017).
9. 2017. Tropo - Cloud API for Voice and Messaging. <https://www.tropo.com/>. (2017).
10. 2017. The World Factbook - Central Intelligence Agency. <https://www.cia.gov/library/publications/the-world-factbook/fields/2103.html#136>. (2017).
11. Sheetal K Agarwal, Arun Kumar, Amit Anil Nanavati, and Nitendra Rajput. 2009. Content creation and dissemination by-and-for users in rural areas. In *Information and Communication Technologies and Development (ICTD), 2009 International Conference on*. IEEE, 56–65.
12. Abhishek Behl and Maneesha Vashishtha. 2014. Innovative IVR system for farmers: enhancing ICT adoption. *International Journal of Information Technology & Computer Sciences Perspectives* 3, 2 (2014), 970.
13. Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement* 20, 1 (1960), 37–46.
14. Brian DeRenzi, Jeremy Wacksman, Nicola Dell, Scott Lee, Neal Lesh, Gaetano Borriello, and Andrew Ellner. 2016. Closing the Feedback Loop: A 12-month Evaluation of ASTA, a Self-Tracking Application for ASHAs. In *Proceedings of the Eighth International Conference on Information and Communication Technologies and Development*. ACM, 22.
15. Mohamed Gulaid and Aditya Vashistha. 2013. Ila Dhageyso: an interactive voice forum to foster transparent governance in Somaliland. In *Proceedings of the Sixth International Conference on Information and Communications Technologies and Development: Notes-Volume 2*. ACM, 41–44.
16. Zahir Koradia, Piyush Aggarwal, Aaditeshwar Seth, and Gaurav Luthra. 2013. Gurgaon idol: a singing competition over community radio and IVRS. In *Proceedings of the 3rd ACM Symposium on Computing for Development*. ACM, 6.
17. Klaus Krippendorff. 2012. *Content analysis: An introduction to its methodology*. Sage.
18. Alexander Lex, Nils Gehlenborg, Hendrik Strobel, Romain Vuillemot, and Hanspeter Pfister. 2014. UpSet: Visualization of Intersecting Sets. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '14)* 20, 12 (2014), 1983–1992.
19. Henry B Mann and Donald R Whitney. 1947. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics* (1947), 50–60.
20. Aparna Moitra, Vishnupriya Das, Gram Vaani, Archana Kumar, and Aaditeshwar Seth. 2016. Design Lessons from Creating a Mobile-based Community Media Platform in Rural India. In *Proceedings of the Eighth International Conference on Information and Communication Technologies and Development*. ACM, 14.
21. Preeti Mudliar, Jonathan Donner, and William Thies. 2012. Emergent practices around CGNet Swara, voice forum for citizen journalism in rural India. In *Proceedings of the Fifth International Conference on Information and Communication Technologies and Development*. ACM, 159–168.
22. Walter J Ong. 2013. *Orality and literacy*. Routledge.
23. N. Patel, D. Chittamuru, A. Jain, P. Dave, and T.S. Parikh. 2010. Avaaj Otalo: a field study of an interactive voice forum for small farmers in rural India. In *Proceedings of the 28th international conference on Human factors in computing systems*. ACM, 733–742.
24. A.A. Raza, C. Milo, G. Alster, J. Sherwani, M. Pervaiz, S. Razaq, U. Saif, and R. Rosenfeld. 2012. Viral Entertainment as a vehicle for disseminating speech-based services to low-literate users. In *International Conference on Information and Communication Technologies and Development (ICTD)*, Vol. 2.
25. Agha Ali Raza, Farhan Ul Haq, Zain Tariq, Mansoor Pervaiz, Samia Razaq, Umar Saif, and Roni Rosenfeld. 2013. Job opportunities through entertainment: Virally spread speech-based services for low-literate users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2803–2812.
26. Robert Roche, Eric Hladilek, and Scott Reid. 2006. Disaster recovery virtual roll call and recovery management system. (April 11 2006). US Patent 7,026,925.
27. Bruce Rocheleau and Liangfu Wu. 2005. E-Government and financial transactions: Potential versus reality. *The Electronic Journal of e-Government* 3, 4 (2005), 219–230.
28. J. Sherwani, N. Ali, S. Mirza, A. Fatma, Y. Memon, M. Karim, R. Tongia, and R. Rosenfeld. 2007. Healthline: Speech-based access to health information by low-literate users. In *International Conference on Information and Communication Technologies and Development, ICTD*. IEEE, 1–9.

29. Aditya Vashistha, Edward Cutrell, Gaetano Borriello, and William Thies. 2015a. Sangeet swara: A community-moderated voice forum in rural india. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 417–426.
30. Aditya Vashistha, Edward Cutrell, Nicola Dell, and Richard Anderson. 2015b. Social media platforms for low-income blind people in india. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility*. ACM, 259–272.
31. Jerome White, Mayuri Duggirala, Krishna Kummamuru, and Saurabh Srivastava. 2012. Designing a voice-based employment exchange for rural India. In *Proceedings of the Fifth International Conference on Information and Communication Technologies and Development*. ACM, 367–373.
32. Nikolas Wolfe, Juneki Hong, Agha Ali Raza, Bhiksha Raj, and Roni Rosenfeld. 2015. Rapid development of public health education systems in low-literacy multilingual environments: Combating ebola through voice messaging. In *ISCA Special Interest Group on Speech and Language Technology in Education (SLaTE)*. INTERSPEECH.
33. Ayesha Zainudeen, Rohan Samarajiva, and Nirmali Sivapragasam. 2010. Cellbazaar, a mobile-based e-marketplace: Success factors and potential for expansion. (2010).