EXPERIENCE WITH DEVELOPING AND DEPLOYING AN AGRICULTURAL INFORMATION SYSTEM USING SPOKEN LANGUAGE TECHNOLOGY IN KENYA

Roger C F Tucker, Outside Echo Ltd
Mucemi Gakuru, Teknobyte Ltd

roger@outsideecho.com, mucemi@teknobyte.co.ke

ABSTRACT

We describe the progress of the Local Language Speech Technology Initiative in Kenya, where since starting in 2003, technology and expertise have been successfully transferred to the Kenyan partners, culminating in the launch of the National Farmers Information Service (NAFIS) in April 2008. NAFIS is primarily a voice service accessed over the phone and offers a wide range of information in Kiswahili or Kenyan English, supplementing the existing agricultural extension services.

Index Terms— Speech Synthesis, Developing Nations, Information Systems

1. INTRODUCTION

Spoken Language Technology (SLT) has great potential for making IT more accessible for those without text or computer literacy, enabling widespread access to information on health, markets, agriculture, weather, microfinance etc. throughout the developing world. However, compared to the use of SLT in the developed world, that potential has yet to be realized. There are two key factors that help explain the difficulty – the lack of success of ICT for Development projects generally, and the specific problem of SLT for local languages in the developing world.

ICT projects in the developing world have had a difficult history, particularly those with a specific developmental aim [1]. The reasons for the many failures are complex but revolve around the issues of appropriate solutions, local ownership and most of all, sustainability.

Spoken Language Technology for the languages of the developing world has also been problematic, due mostly to the lack of speech data and linguistic resources, the skills shortage in the countries where the languages are spoken, and the general lack of funding for research in the poorer countries of the world.

In this paper we describe how we have addressed both these factors in the Local Language Speech Technology Initiative (LLSTI), and the progress made by the Kenyan LLSTI partners towards the development and deployment of an agricultural information system using SLT. For simplicity we will use the terminology "Northern" for the developed world and "Southern" for the developing world in the text.

2. LLSTI

The Local Language Speech Technology Initiative [2] was set up in 2003 as a partnership between a number of Southern and Northern institutions, with the aim of creating communities of interest in SLT for Development throughout the global South. LLSTI was initially led by Outside Echo with financial support from the Department for International Development in the UK and the International Development Research Centre in Canada, but has since progressed to an informal network of partnerships.

Until LLSTI started, SLT for Southern languages were most often developed in a Northern lab, with a speaker of the language temporarily joining a team of experts to annotate data and advise on language characteristics. This approach has a number of drawbacks, not least that SLT requires on-going development and often needs tailoring to a specific application. Unless the team lives and works in the country where the systems will be deployed, it is difficult to do this effectively.

A key issue which the LLSTI approach addresses is that of technology transfer and ownership. By allowing the Southern partners to take ownership of the technology development process early on, they have the motivation and maturity to see it through to successful deployment.

2.1. LLSTI in Kenya

The University of Nairobi was approached early on in the LLSTI project to undertake the development of a Kiswahili Text to Speech (TTS) system using Festival, which has been reported on elsewhere [3]. Sponsorship for the development came from Oneworld, who in turn had received funding

from the Vodaphone Foundation to set up a mobile-based social enterprise in Kenya, which subsequently was launched as Mobile4Good. The initial plan had been for Mobile4Good to incorporate voice into their SMS-based business model, but in the end there was simply not enough financial resource to make this happen.

At the end of the first phase of development, the Kiswahili system was taken up by a Kenyan company Teknobyte and further improved. In addition, a demonstration Kenyan English system was developed from an open-source US English TSS by simply rerecording the US English unit-selection database and performing automatic labeling using the Hidden Markov Model Toolkit (HTK).

At this point (mid-2005), we searched for some sponsorship to re-engineer the TTS to production standards and optimize the execution speed, so it could be made widely available to application designers, and also used in screen-readers for the blind. A proposal for the latter in conjunction with the Kenyan Society for the Blind has always looked very promising, but it is taking an impossibly long time for the funding to come through.

Apart from re-engineering, the next step was to build a pilot application using the Kiswahili TTS which would demonstrate to the development community the potential of a voice-based information system. Kiswahili is the first language of almost 800,000 people and over 100 million rural people in Eastern Africa are second-language users

3. AGRICULTURAL INFORMATION PILOT

We identified farming as the sector where it would be easiest to make a direct impact on Kenyan society, since as much as 75% of the population is involved in farming in some way. A recent UN survey of e-agriculture [4], identified information exchange and communication processes as critical, highlighting among other things the following areas:
- Enhancing farmers and producers access to markets and information on farming techniques & practices
- Improving dissemination of and access to scientific and technical information

We worked in partnership with the National Agriculture and Livestock Extension Programme (NALEP) to provide a phone-based information system for growing bananas. Bananas were chosen because of the demand for knowledge of this particular crop, and the system covers all aspects from deciding whether or not to grow bananas through to harvesting, pest control and market information. In addition bananas are grown in all arable parts of the country, which meant that the system would find use in many parts of the country thereby giving a chance for testing various types of data.

At first we had hoped to engage Safaricom (the Kenyan Vodaphone subsidiary) in the deployment of the pilot. Previous meetings with Safaricom had been very encouraging, even with a promise of looking for a suitable pilot application themselves. However, by the time the Banana line was under development there was some scepticism in the company over development-focused applications. An SMS service to deliver market price information (the most frequently cited information need in development literature!) was proving unsuccessful, with very few subscribers indeed. The additional running cost of a voice-based service did not make much sense in the light of this poor take-up of the SMS service. Because of Safaricom's cool response, we partnered with a small IVR company, SpeechNet, to host the Banana Line.

The content was created in XML and XLST transforms were used to create
- HTML (for visual checking)
- Festival scripts to synthesize the text
- Scripts to drive the Ubuntu-based IVR

The content originated in English and was translated into Kiswahili by hand. On phoning the system, the user is given the choice of either language. The English TTS used was a British English female voice (Nina) – the Kenyan English system was too buggy to be used at this stage.

The pilot [5] included a formal test on a carefully selected group of 10 farmers in the Kirinyaga district of Kenya. 7 out of the 10 chose to listen in English, but then struggled with the British accent. Those that chose Kiswahili loved the voice, but then struggled with the formal Kiswahili grammar used in the translation. This could be attributed to the fact that many had not learnt the language formally in school and Kiswahili grammar as spoken has changed to accommodate borrowed words as well as phrases from English and local languages. All said they liked the voice system and would prefer it to written material, but it was clear that the accent and translation issues would need to be fixed before a real deployment was made. Subsequent implementations such as in Nafis, described below, have tried to use borrowed words and phrases as currently spoken in Kiswahili. The farmers were also played some speech from the demonstration Kenyan English TTS system, and found it even clearer than the Kiswahili.

Consulting with NALEP after the pilot (late 2006) gave a comprehensive picture of what a real service could offer:
- Different crops with separate phone numbers
- Simple update of information
  - Website where anyone can update or enter new information, immediately becomes available in the telephone system

- o Updating at district/provincial level by extension workers
- o Farmers can also add information or post questions to the system
- Up-to-date information
  - o Commodity prices & other market information
  - o Weather
  - o Urgent announcements e.g. disease outbreaks
- Local information
  - o Local market information
  - o Suitable crops to grow in a specific area
  - o Contact information to the local agricultural extension officers
- Personal information
  - o Enter acres of land -> get calculations of possible revenue, required investments etc.
  - o Enter time of planting -> get instructions suitable for the stage your plants are in
  - o Caller-id and pin-code for ID so info doesn't have to be re-entered.
- Option in the voice system to get key information sent by SMS or e-mail

After the pilot, Outside Echo (the UK partner) was able to pull back and the local LLSTI partner, Teknobyte, took over full leadership of the ongoing developments.

4. KENYAN ENGLISH TTS

A key result from the Banana line pilot has been the need for a Kenyan English voice. It is easy to undervalue the importance of a voice in the local accent, even if the language itself is a world language. Most of the English TTS voices available have American or British accents. Even though such voices are fairly intelligible to many Kenyans, a greater majority often face difficulties in picking up entire phrases due to differences in intonation and pronunciation, especially of the large number of borrowed words from local languages. The development of a TTS based on the Kenyan English dialect was therefore deemed necessary in order to produce a system whose output would be more natural to the Kenyan population and hence more easily understood.

Teknobyte set out to develop the Kenyan English by first defining a phoneset to capture all the sounds in the dialect including those of the borrowed words. A lexicon of 10,000 most commonly used words in the Kenyan English was then developed based on this starting phoneset. A trial Kenyan English voice was then developed using this lexicon and the recording from the demonstration Kenyan voice described earlier, in order to find out whether the lexicon would yield an improved voice. Adjustments were carried out on the phoneset and the corresponding lexicon repeatedly until we were satisfied that they captured the Kenyan English dialect as best as possible. Having settled on a final phoneset, the lexicon was built up to 60,000 words including 6,000 borrowed words. Every so often, as the lexicon continued to build up, a new voice would be made and it was observed that the voices continued to improve.

A speech database was then developed by first collecting a large text corpus comprising of 20,000 sentences, from various sources such as novels, newspaper articles and written speeches. The novels used were those that were rich in local names. The corpus was then checked for grammatical and typographical errors, like improper punctuations and spelling mistakes, as this would impact on the transcription. A tool was then developed to arrange the corpus into a list of sentences and then transcribe them. The transcribed sentences were processed further to define phone-position using specific symbols that indicated the sentence, phrase and word boundaries.

Selection of sentences for the speech database was done based on units [6] such as phone-phone combinations to obtain a 1500-sentence phonetically balanced database. A professional Kenyan newscaster, Frank Muiruri, was selected to record the 1500 sentences, as he was aware of the language features and also well trained in use of microphones. This process yielded a good Kenyan English voice which has been used in development of NAFIS as described below.

5. NATIONAL FARMERS INFORMATION SERVICE (NAFIS)

Due to the importance of Agriculture in Kenya's economy, which is mostly done by small scale farmers, output can be improved by enhancing technical support through extension services. The number of small scale farmers has continued to rise while that of extension officers has remained largely unchanged owing to the limited resources afforded by the government budget. This has therefore created the need to harness ICT Technology as a cost effective supplement to enhance the reach of extension services to the 4.5 million farmers. As a result the idea to develop NAFIS was conceived. It is primarily a voice service (automated IVR) for providing agricultural extension information through telephony, offering the user a choice of either Kenyan English or Kiswahili. It builds and improves on the agricultural pilot described in Section 3 by including more enterprises as well as features, and uses the Kenyan English TTS system described above.

The information obtained through the phone is deliberately summarised to save on the time the farmer has to spend on the call. Furthermore, the system is interconnected to all the Telephone service providers, thus allowing cheap calling rates by farmers as they do not pay interconnection charges. NAFIS is updated through the Web so as to enable updates

by field extension officers, with the information so updated being accessed through the phone, while being available on the website. When fully developed, this very crucial feature of NAFIS will allow farmers to access regional based information which is relevant to their agro-ecological zones.

NAFIS provides the following benefits:
- Offers critical regional based information to farmers at the appropriate time
- Creates demand for extension services by informing farmers where to get to their extension officer.
- Creates a ready communications channel between the providers of extension services and the farmers
- Enables the government to get feedback on services needed most by farmers

The most tangible benefit of NAFIS so far has been the feedback it has given the Ministry of Agriculture on the amount of information farmers are currently receiving and what their information needs actually are.

Clearly, NAFIS does not substitute the extension officer or other existing sources of information, but rather enhances them. It is noted here that crucially, NAFIS fills in a gap in the existing information chain. In fact the experience with developing and deploying NAFIS has led to the realisation that to provide an effective information delivery mechanism, many platforms need to be carefully integrated in such a way as to harness the benefits of individual platforms while avoiding the limitations of each one of them. In other words, an effective solution must contain as many of these individual platforms as possible organised carefully so that each falls in its place to fulfill a set objective. The solution then becomes an information matrix whose elements are the individual platforms and whose composition is done on the basis of the set objectives.

Further development and deployment of NAFIS will therefore be done in the context of an agricultural information matrix that will comprise among others:
- A fully developed web-portal with illustrative pictures
- A complete feedback mechanism from all the sources of information
- Integrating a Short Message System
- An efficient content gathering mechanism
- A well worked out awareness creation campaign

6. CONCLUSION

The development of an information system using SLT for Kenya has already taken five years and is not yet fully complete. Much of the elapsed time has not been development time, but simply the time needed to grow partnerships and for stakeholders to take on board the new possibilities and think them through properly.

This seems slow compared to the alternative approach of a large donor or government department to have simply bought lots of Northern resource, put together the system, installed it, and handed it over to NALEP. However, the experience of many years of ICT for Development suggests that this approach is unlikely to have worked. Instead, we have local expertise and ownership, and every possibility for a successful, locally maintained, appropriate information system for Africa.

7. ACKNOWLEDGEMENTS

8. REFERENCES

[1] R Heeks, "Failure, Success and Improvisation of Information Systems Projects in Developing Countries", 2002. Available from: *www.sed.manchester.ac.uk/idpm/ research/publications/wp/di/*

[2] R Tucker and K Shalonova, "Supporting the Creation of TTS for Local Language Voice Information Systems", Interspeech 2005, Lisbon, Portugal, pp. 453-456, 2005

[3] M Gakuru et al, "Development of a Kiswahili Text To Speech System", Interspeech 2005, Lisbon, Portugal, pp. 1481-1484, 2005

[4] International Telecoms Union, *e-Agriculture Survey Analysis, 2007*. Available from: *www.itu.int/wsis/c7/e-agriculture/docs/survey-analysis-2007.pdf*

[5] P Näsfors. "Efficient Voice Information Services for Developing Countries", Masters Thesis, Department of Information Technology, Uppsala University, Sweden, 2007. Also available from *www.llsti.org/documents.htm*

[6] M Gakuru et al, "Design of Speech Data Base for Unit-Selection in Kiswahili TTS", E-Tech2004, Kenya. Also available from *www.llsti.org/documents.htm*