

ru | STEVIN

Dutch-Flemish research programme for
Dutch Language and Speech Technology

STEVIN mid-term Internal Evaluation



May, 2008

Introduction

This (self-)assessment of the STEVIN programme by the STEVIN programme Committee has been organized around nine questions. A separate section is dedicated to each of these questions and their answers.

These nine questions are:

1. Do you think the STEVIN programme is sufficiently reaching its main aims and objectives described in the programme text? That is:
 - Realise an effective digital language infrastructure for Dutch, based on the BATAVO priorities (i.e., 'BASisTaal&spraakVOorzieningen' = BLARK 'Basic LAnguage Resources Kit');
 - Carry out strategic research in the field of language and speech technology, especially in areas of high demand for specific applications and technologies;
 - Advance the creation of networks and the consolidation of language and speech technology activities, educate new experts, and promote discussion and transfer of knowledge.
2. Does the selection of the research projects reflect the stratified chain approach of the STEVIN programme, i.e. address all four levels in an appropriate way?:
 - The basic language resources, explicitly limited to the data and their production;
 - HLT research and development, ultimately resulting in available HLT components;
 - Application development (HLT embedding);
 - The user ('vraagstimulering')
 - Are there any critical gaps in the programme that should still be addressed?
3. Is the (scientific) output of the STEVIN projects adequate in terms of quality and quantity? Which are the most outstanding results? Are the results made available in a non-discriminative way? Do you think that IPR policy is adequately dealt with in the STEVIN programme?
4. One aim of the programme is to stimulate collaboration between academia and industry in Flanders and the Netherlands. Is there enough evidence indicating that this aim is being achieved?
5. An integral part of the STEVIN programme is its supporting programme which specifically aims at organising and/or financially supporting activities that stimulate network relations between the different types of actors that are involved in developing, implementing or embedding HLT technologies and encourage knowledge transfer between these actors. How do you judge the success of this part of the programme so far?
6. Are the STEVIN programme and its results visible to the (inter)national field of HLT research, the interested industry, and the interested public? Please provide some examples.
7. What is the (scientific) impact of the STEVIN programme with respect to Dutch and Flemish HLT?
8. Is the organisational structure of the programme, namely the International Assessment Panel, HLT Board, programme committee, and programme bureau appropriate? Does the programme offer sufficient guarantees on transparent, impartial and objective proposal evaluation, project monitoring and decision taking in general? Are potential conflicts of interest adequately dealt with on all levels of the programme?
9. Were the installed procedures for granting, launching and monitoring the projects and for transferring and approving the project results adequate?"

Before turning to these questions, however, we give an overview of the calls launched under the STEVIN programme, since there will be multiple references to these calls in the sections to follow.

The STEVIN funding instruments and their respective budgets were:

- I) three open calls for R&D proposals - M€ 7
- II) three tenders - M€ 1,65
- III) three calls for demonstration projects - M€ 1
- IV) three calls for educational projects - M€ 0.2
- V) a continuously open call for networking grant - M€ 0.1

More information about the calls and legal foundations for these calls are incorporated in the STEVIN Fact File and so is a listing of the assessment criteria used and a description of the assessment procedures.

Question 1

Do you think the STEVIN programme is sufficiently reaching its main aims and objectives described in the programme text? That is:

- *Realise an effective digital language infrastructure for Dutch, based on the BATAVO priorities (i.e., 'BASisTaal&spraakVOorzieningen' = BLARK 'Basic LAnguage Resources Kit');*
- *Carry out strategic research in the field of language and speech technology, especially in areas of high demand for specific applications and technologies;*
- *Advance the creation of networks and the consolidation of language and speech technology activities, educate new experts, and promote discussion and transfer of knowledge.*

The main target of the STEVIN programme is threefold:

1. *The realisation of an effective digital language infrastructure for Dutch, based on the BATAVO priorities,*
2. *The execution of strategic research in the field of language and speech technology*
3. *The creation of networks and the consolidation of language and speech technology activities, educate new experts, and promote discussion and transfer of knowledge*

STEVIN objective 1: Realisation of an effective digital language infrastructure for Dutch

The main aims of the STEVIN programme were laid down in the STEVIN Programme description¹ which was mainly based on an earlier study commissioned in 2002 by the Dutch and Flemish government in which the main items of the basic language resources kit (BLARK) were listed that needed to be developed in order to strengthen the economical and cultural position of the Dutch language in the modern ICT-based society.² In that report it was stressed that an effective digital infrastructure of Dutch Language resources is essential in language and speech technology for both the research and the industrial community.

A 'digital infrastructure for Dutch Language Resources' involves at least three aspects:

1. Relevant language resources must exist or be created and guidelines, conventions, and best practices for creating these and other language resources must be made.
2. The relevant language resources must be stored and maintained, and these resources with their documentation, as well as guidelines, conventions and best practices must be made easily available to all researchers and developers in a non-discriminative way
3. The intellectual property rights (IPR) surrounding the language resources must have been adequately dealt with, so that the resources can actually be used by researchers and developers.

We discuss each of these aspects in turn.

Language Resources

In the STEVIN programme three open calls for research and development (R&D) projects have been launched. Because the availability of adequate resources is often a necessity for basic as well as for more applied research, several of the projects granted in the first call for proposals were focused on the collection and elaboration of resources (both corpora and tools), and guidelines for creating them. Examples are D-COI and Jasmin-CGN. D-COI was meant as a pilot study for one of the main goals of the STEVIN programme: the creation of a huge (>500M words) Corpus of Written Dutch. Also several projects from other calls were focused on creating new or extending and enhancing existing language resources (e.g. DPC, Lassy, STEVIN can PRAAT). Projects (from all calls) were always stimulated to yield language resources as a side benefit next to their their main research objectives, and many awarded projects have done so (examples are Autonomata, IRME, COREA, Daeso).

¹ [STEVIN Meerjarenprogramma \(in Dutch\):](http://taalunieversum.org/taal/technologie/stevin/documenten/stevin_mjp_150904.pdf)

http://taalunieversum.org/taal/technologie/stevin/documenten/stevin_mjp_150904.pdf

² [BaTaVo-prioriteitenlijst \(in Dutch\):](http://taalunieversum.org/taal/technologie/stevin/documenten/batavo.pdf)

<http://taalunieversum.org/taal/technologie/stevin/documenten/batavo.pdf>

In the three open calls the selection of project topics was left to the researchers and developers. Though the projects submitted were of course ranked based (inter alia) on conformity to the STEVIN programme goals, this did not guarantee that all or even the most important goals would be realized in the STEVIN programme. In order to ensure the creation of a number of high priority language resources, three specific tenders have been issued:

1. one for the creation of a semantic lexicon, resulting in the project Cornetto
2. one for the creation of a state-of-the-art Dutch Speech Recognizer and speech recognizer toolkit, resulting in the project SPRAAK
3. one for a richly annotated corpus of written Dutch containing more than 500M words from a wide variety of sources, resulting in the project proposal SONAR

The SPRAAK and Cornetto projects are in their final phases or have finished already.

The realization of the Written Corpus tender had been partially prepared and foreshadowed in previous projects of round 1 and 2 (D-COI, LASSY, COREA, ...). Due to a variety of reasons that will be described below (in the section on Question 8), the SONAR project granting has been delayed, so the bulk of the work for this project still has to be carried out.

Despite this delay, it can reasonably be expected that the development of a large, annotated corpus of written Dutch will take place. This is important since such a written corpus was one of the main targets of the programme as a top priority both for the language and speech research communities and as the top priority in the BaTaVo report as well. The results of the D-COI project, a proposal granted in the first call, which was set up as a pilot study for the final corpus and which collected and processed some data (50M) and did a study how to compile and annotate the final, full-scale corpus of written Dutch, can now be fully exploited.

Availability, Accessibility, and Maintenance

In order to guarantee the availability of the language resources, both resources that existed prior to the STEVIN programme as well as resources that have been or are being created in the STEVIN programme are being made available via the Dutch HLT Agency (TST-Centrale). The Dutch HLT Agency has been started up in 2003, to serve as a central portal and service centre for Dutch language and speech technology resources.

The main tasks of the Dutch HLT Agency are acquisition and IPR, management and maintenance, and distribution and service. A pricing committee has been formed to support the HLT Agency in pricing matters. The current pricing policy distinguishes (non-commercial) research and commercial use. Commercial use prices are in accordance with the market, for (non-commercial) research resources are freely available.

IPR

Ensuring the correct treatment of intellectual property rights (IPR) is a non-trivial issue, especially if it has to be done by researchers and developers who are not particularly knowledgeable in this area. The newly created Dutch HLT Agency attempts to support the researchers, but also for them most of the matters relating to this aspect are new. This caused problems and delays for several of the projects, especially (but not exclusively) from the first call. For this reason, the programme committee proposed to set up an IPR committee, with members with some experience in this field from the programme committee (including the Dutch HLT Agency representative) and HLT Board, supported by a lawyer with expertise in this area. This committee has given advice to several projects (for example D-COI, DPC, SPRAAK) and produced several other results, in particular a leaflet (c.f. STEVIN Fact File) to help convince external data providers (e.g. publishers) to make available their data and give them the confidence that doing this is not harmful to them or hurting their normal business interests, a schematic overview of the relations between actors involved (c.f. STEVIN Fact File) and a number of standard contract templates³ that can be used by project responsables as a

³ [standard contract templates](#):

http://www.inl.nl/index.php?option=com_content&task=blogcategory&id=103&Itemid=5

basis for dealing with IPR issues in an adequate manner. In addition, it has been decided that an Open Source strategy towards the IPR issues is a possibility within STEVIN, and some projects have delivered or will deliver their results in this way (e.g. SPRAAK). These results have been very useful, but they do not imply that all IPR issues are now solved, and it is certainly true that IPR issues have led to delays in delivering some of the project results.

Guidelines, conventions and best practices have been created in several projects, and are also available via the HLT Agency. They have of course also been made available in the form of scientific publications, e.g. in LREC conferences and proceedings.

STEVIN objective 2: The execution of strategic research in the field of language and speech technology

The first call already resulted in some projects carrying out strategic research, including COREA, IRME, and Autonomata.

The second call for proposals was mainly focused on strategic scientific research as can be seen in the overview of the granted projects in the STEVIN Fact File. The granted projects include DAESO (Detecting and Exploiting Semantic Overlap), Midas (Missing Data Solutions), and LASSY (large scale syntactic annotation and research into improving a robust syntactic parser for Dutch).

The third and last call for R&D proposals was focused on strategic applied research, using where possible the results of the previous projects. All five granted projects involve applied research.

Together with the demonstration projects (see below) the STEVIN programme shows a clear direction from resources to real-world demonstrators. All strata in the stratified chain approach are covered in the STEVIN programme.

STEVIN objective 3: The creation of networks and the consolidation of language and speech technology activities, educate new experts, and promote discussion and transfer of knowledge

In order to achieve the third goal, a series of supporting activities, using $\pm 10\%$ of the total STEVIN budget have been planned and have in part already been carried out. These activities include calls for demonstration projects, calls for educational projects, funding of networking activities (such as conferences, workshops, etc.), the organisation of brokerage events related to STEVIN calls, the organization of internal STEVIN Programme Meetings in which STEVIN-funded researchers and developers can discuss their projects and results.

More details about these activities, which we judge to be very successful, will be provided in the section on Question 5.

Summary

In summary, given the initial target of the STEVIN programme ("effective digital language infrastructure", "strategic research in the field of LST" and "creation of LST-networks, education of new experts, and promote discussion and transfer of knowledge") many targets already have been reached in this stage and all targets can reasonably be expected to be reached: Most of the other high-priority goals were or will be handled by the granted proposals.

Tables with the granted projects, classified by type, and the priorities they address can be found in the STEVIN Fact File.

Question 2

Does the selection of the research projects reflect the stratified chain approach of the STEVIN programme, i.e. address all four levels in an appropriate way?:

- The basic language resources, explicitly limited to the data and their production;
- HLT research and development, ultimately resulting in available HLT components;
- Application development (HLT embedding);
- The user (‘vraagstimulering’)
- Are there any critical gaps in the programme that should still be addressed?

The STEVIN Programme started with the ambition to fund four types of R&D activities:

- I. Development of corpora and language resources for the Dutch language as written and spoken in the Netherlands and Flanders (including tools for annotating and processing corpus data)
- II. Strategic research aimed at the development of novel procedures and algorithms for speech and natural language processing, with a strong emphasis on processing Dutch
- III. Application oriented research, aimed at creative deployment of existing resources and tools for novel applications, potentially combined with the enhancement of the functionality and/ or accuracy of existing techniques and tools
- IV. Development of operational applications aimed at propagating the usefulness of Language and Speech Technology for economically and socially relevant problems.

As to the first objective, the original priorities in the domain of language technology were:

1. A large corpus of written Dutch
2. An electronic lexicon
3. Parallel corpora

Projects which aim at the construction of these resources are:

- D-COI: preparation for the construction of a 500 million word corpus of written Dutch.
- LASSY: a 1 million-word tree bank of written Dutch, based on D-COI.
- CORNETTO: a lexical semantic database for Dutch, covering 40 thousand entries.
- DPC: creation of a 10 million word sentence-aligned parallel corpus for Dutch-English and Dutch-French.

Besides, there are projects which have the creation of annotated corpora or lexical resources as one of their secondary objectives:

- COREA: a limited amount of data annotated with co-referential information
- IRME: a lexical database of Dutch multi-word expressions, covering 5 thousand entries
- DAESO: an annotated monolingual Dutch parallel corpus of 1 million words

The original priorities in the field of speech technology were:

1. speech and multimodal corpora for:
 - o applications such as CALL (Computer Assisted Language Learning);
 - o applications in which names and addresses play an important role;
 - o CCQA applications (questions and answers in call centres), educational applications;
2. multimodal corpora for applications of broadcast news transcription or person identification;
3. text corpora for the development of stochastic language models;
4. tools and data for the development of:
 - o robust speech recognition;
 - o automatic annotation of corpora;
 - o speech synthesis;

Projects which address these objectives are:

- AUTONOMATA: a corpus of spoken person and address name utterances. Since many non-natives and young people are recorded, this corpus may also be relevant for CALL and educational applications for children.

- JASMIN-CGN: an extension of the Spoken Dutch Corpus with speech of children, non-natives, elderly people and human-machine interaction. It can therefore be seen as support for applications such as CALL and educational applications for children.
- D-COI and SONAR: for the construction of a large Dutch written corpus, which can, inter alia, be used for stochastic language modelling.
- SPRAAK: a speech recognizer toolkit and speech recognizer for Dutch
- STEVIN can PRAAT: extension of the functionalities of the PRAAT program

All bullets of the first objective for speech technologies can be said to have been satisfactorily addressed, with the exception of the second bullet: multimedia corpora have not been produced in the STEVIN programme and are not planned to be produced. There also were no real projects in the area of speech synthesis, in part because commercial speech synthesis systems are available for Dutch and of high quality. However, one can reasonably state that Autonomata can make a small contribution to improvements of the synthesis of names.

Overall it can be stated that the first objective has been addressed almost completely and it has to a large extent already been achieved (5 out of the 6 projects have finished already or are in their final stages).

Turning to the second objective, the priorities in the domain of language technology were:

1. semantic analysis (tagging, integration with syntax and morphology)
2. text pre-processing (tokenization, spelling correction, named entity recognition, ...)
3. morphological analysis (compounding and derivation)
4. syntactic analysis: a robust parser for Dutch

Projects which primarily address these priorities are:

- COREA: development of a robust system for automatic co-reference resolution.
- IRME: methods and tools for the automatic identification and representation of multi-word expressions.
- DAESO: tools for automatic alignment, classification of semantic relations and text-to-text generation

These projects cover aspects of semantic analysis. Parts of D-COI cover other aspects of semantic analysis (thematic roles and spatio-temporal relations), as well as various aspects of text pre-processing. LASSY combines the construction of a treebank with the further development and refinement of a robust parser for Dutch. A major lacuna is the absence of a project on morphological analysis.

The original priorities in the field of speech technology were:

1. robustness of speech recognition;
2. output treatment (inverse text normalization);
3. confidence measures;
4. adaptation;
5. lattices.

They are addressed by:

- AUTONOMATA: grapheme to phoneme conversion for Dutch and Flemish names.
- MIDAS: noise reduction techniques and confidence measures, to be integrated in SPRAAK
- N-BEST: benchmarks for the evaluation of speech recognizers for northern and southern Dutch
- SPRAAK: a speech recognizer toolkit and speech recognizer for Dutch

So, also for the second objective it can be stated that it has almost completely been addressed. Two out of four projects have already finished, so also the realization of this objective is well underway.

As for the **third** objective, the priorities for both speech and language technology were:

1. Information extraction (monolingual and multilingual)
2. Semantic Web
3. Dialogue systems and Question-Answering
4. Summarization and text generation
5. Machine translation
6. Teaching, learning, gaming

Projects directly addressing these objectives are:

- PACO-MT: machine translation for Dutch, English, French.
- DuOMan: resources and tools for identifying and aggregating sentiments in online data sources.
- DAISY: summarization technology.
- DISCO: development and integration of speech technology into courseware for language learning

Besides, many of the projects mentioned under objective 2 have the integration in applications as a secondary objective: AUTONOMATA aims at ASR and TTS, MIDAS at ASR, and COREA and DAESO aim both at information extraction, summarization and Question-Answering.

All six of the application areas are covered, except perhaps the Semantic Web.

As to the **fourth** objective, it is mainly addressed by the demonstrator projects. These are discussed in the section on question 5.

Summary

The following table provides the percentages of STEVIN funding per stratum and by language or speech technology. One can observe that all strata in the stratified layer approach have been covered. The focus has clearly been on the creation of resources, but strategic research, applied research, and technology demonstration projects aimed at the user are also well represented.

Percentage of STEVIN funding per STEVIN priority		
• Speech technology resources	21,5%	
• Language technology resources	29,0%	
% STEVIN funding for basic resources		50,5%
• Speech technology research	14,5%	
• Language technology research	9,1%	
% STEVIN funding for basic research		23,6%
• Speech technology application-oriented research	7,4%	
• Language technology application-oriented research	8,2%	
% HLT Application-oriented research		15,6%
• Speech technology demonstration projects	3,7%	
• Language technology demonstration projects	6,6%	
% HLT Demonstration projects		10,3%
% STEVIN funding for speech technology		47,1%
% STEVIN funding for language technology		52,9%

The following picture illustrates how the project types (classified by stratum) have been distributed over time in the STEVIN programme:

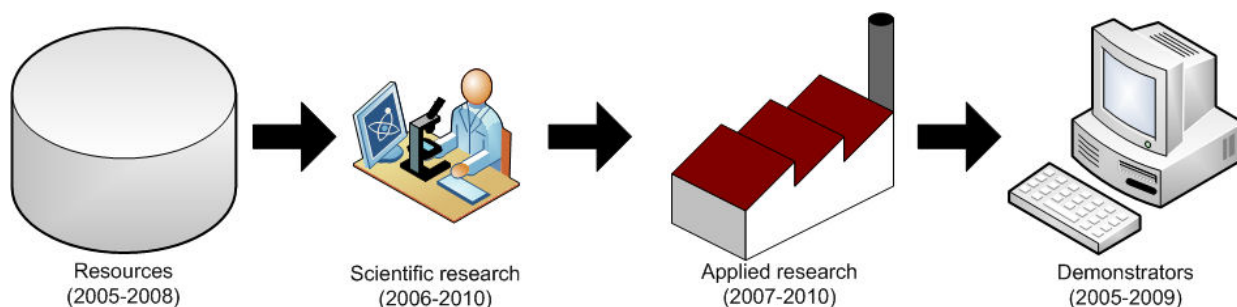


Figure 1: Overview of the focus and time span of the various calls in the STEVIN programme

Question 3

Is the (scientific) output of the STEVIN projects adequate in terms of quality and quantity? Which are the most outstanding results? Are the results made available in a non-discriminative way? Do you think that IPR policy is adequately dealt with in the STEVIN programme?

Outputs of the STEVIN project

The projects that were granted in the first and second call of the program have thus far (mid January 2008) produced 15 scientific papers in international conferences such as LREC (Language Resources and Evaluation Conference), Interspeech, CLI (Congress of Linguistics), etc., 13 papers in proceedings of international workshops (at ACL, COLING, IJCNLP, etc.), and 2 book chapters. Several papers are currently under review. As can be derived from the bibliography included in the STEVIN Fact File, the majority of published papers also describe collaborative research.

Until now five projects have delivered, or are in the process of doing so, their results to the HLT Agency. The most outstanding results are listed below (ordered per project and with the projects listed in alphabetical order):

- [AUTONOMATA] A tool suite for pronunciation lexicon development: the tool is built around the state-of-the-art Nuance g2p converter, but it also contains post-processors for improving the g2p conversion in special domains such as person names and geographical names, as well as tools for creating similar post-processors for other domains.
- [AUTONOMATA] A collection of 72K spoken utterances of command words (16% of the utterances), person names (42% of the utterances) and geographical names (42% of the utterances), each supplied with an orthographic and an auditory verified phonetic transcription. The names and the speakers are balanced according to language: 50% native and 50% non-native (equally divided over three languages). This corpus is intended to become a valuable resource for research towards better cross-lingual speech recognition.
- [CGN-JASMIN] A collection of 110 hours of speech recordings representing an extension of the Spoken Dutch Corpus (CGN) along three dimensions: (1) age, because speech of children and elderly people were recorded, (2) mother tongue, because non-natives were recorded, and (3) communication setting, because 50% of the speech material was collected in a human-machine interaction setting that was not envisaged in the CGN (the remaining 50% was read speech). All recordings are delivered with a verbatim orthographic transcription, a transcription of the human-machine interaction (HMI) phenomena, POS tagging and an automatic phonetic transcription.
- [COREA] A significant collection of text material (about 150K words) that was supplemented with co-reference relations (ident, pred, bridge and quant), and that is made available together with a tool for the dynamic visualization of these relations in a standard web browser.

- [COREA] A state-of-the-art co-reference resolution system for Dutch that can be made available as a web-service for document annotation.
- [D-COI] A pilot text corpus of 54MW that is automatically tagged, lemmatized and syntactically parsed. The POS tagging and lemmatization has been verified for 500KW, the syntactic annotation for 200KW. A small part of the data has also been enriched with semantic annotations (semantic roles and temporal adverbial modifiers). All data can be exploited using the adapted version of COREX, the software that was developed for the exploitation of the CGN.
- [D-COI] A set of annotation protocols, corpus development tools and technical reports that can/will be taken as a basis for the production of the final Written Dutch Corpus.
- [IRME] A documented lexical database, called DuELME, containing 5000 manually verified Dutch Multiword entries. The database comes with a graphical user interface to access the data and especially to extend or enhance the database.

Availability of the outputs

Almost all delivered speech and language resources will become available in a non-discriminative way for scientific as well as for commercial utilization. The pronunciation lexicon development tool suite is an exception, because it is a resource that was developed with the aid of background knowledge (lexical data) coming from industries. Therefore, there are justifiable restrictions (laid down in an agreement with the Dutch HLT Agency) to the granting of a commercial license to direct competitors of these industries (defined as providers of the same type of technology, as distinct from application developers aiming at using the technology).

The tool suite for pronunciation lexicon development is also an example of a resource containing a commercial component (the Nuance g2p converter). To run the tool suite, one has to acquire a license for this component, but this license comes cheap for scientific use. A few other resources (e.g. results of CORNETTO) have also been developed with background knowledge coming from industries. As a result, a base fee for commercial use has been agreed.

IPR policy

It has proven difficult to establish an IPR strategy for text corpus distribution that can convince publishers to supply material that will after all be disclosed to users of that corpus (mostly researchers and students). One has finally been able to conceive a model license agreement that has already been approved by some of the biggest candidate text data providers, but the IPR issues have severely hindered a quick start of the D-COI project and are still playing a part in the creation of the envisaged large Written Dutch Corpus.

It has been difficult to develop proper IPR strategies for the distribution of software, because it involves a number of possibly conflicting issues. The aims of the STEVIN program is to make the software available in a non-discriminative way but without hindering a commercial use of this software nor discouraging industries to contribute with background knowledge to the creation of high quality software. Furthermore, the distribution should be organized in such a way that the users know who to contact in case of deficiencies, and that they also know whether bugs can be fixed or not. Although open source models (there are many) can obviously achieve good non-discriminative availability, they are not always solving the other issues. The current viewpoint is not to enforce one single distribution model for all software products, but to search in each individual case for the model that maximally realizes the envisaged objectives. The hope is of course that we will soon have reached a stage in which formerly negotiated distribution models can be reused for products delivered by projects of the 2nd and 3rd call. Although the development of software distribution models has taken a long time, and has delayed the publication of the pronunciation lexicon development tools delivered by AUTONOMATA, it will most probably not delay the publication of future products anymore. The Speech Recognition Toolkit that is scheduled for mid 2008 will hopefully prove this case.

Question 4

One aim of the programme is to stimulate collaboration between academia and industry in Flanders and the Netherlands. Is there enough evidence indicating that this aim is being achieved?

The STEVIN program is achieving a considerable amount of collaboration between academia and industry in Flanders and the Netherlands.

This is clearly demonstrated in a review of the composition of the consortia that are submitting and executing the STEVIN projects: 8 out of the 11 ongoing projects from the two open calls are executed by a consortium that includes one or several industrial partners, while they have a focus on fundamental research and are by definition coordinated by an academic institution. Also 1 of the 2 running projects that are the result of a tender call involves an industry partner. The demonstration projects show less evidence of collaboration: here the participants are primarily commercial SMEs, and only 6 of the 14 projects involve an academic institution - all from the second and third calls (an indication that the active stimulation of collaboration by STEVIN is beginning to bear fruit?). In total, more than 30 different companies are directly involved in STEVIN projects, 23% from Flanders, 77% from the Netherlands, quite a significant portion of the local language and speech industry.

Under the heading "Flankerende Activiteiten" ("supporting activities"), STEVIN organizes various events and activities that explicitly aim to further promote cross-fertilization between industry and academia. The various brokerage events (e.g. "Taal in bedrijf" - 'language@work') attract a significant amount of attention from both parties, and lead to informal networking and formal collaboration between them. The first "Taal in Bedrijf" was attended by 290 participants (260 Dutch/30 Flemish). Many participants had an industrial background (183) but also academic representatives attended (83) and representatives from government agencies (24). Industrial participation was specifically boosted by allowing HLT companies themselves to directly invite their relations and customers to attend the congress.

Based on the findings of the "nulmeting", which records the state of affairs before STEVIN, there already was a high degree of collaboration between academia and the language and speech technology industry, due to the fact that the companies involved show an unusually high amount of investment in research. This naturally promotes collaboration and cross-fertilization, but it also justifies the need for giving this aspect explicit attention within STEVIN.

Question 5

An integral part of the STEVIN programme is its supporting programme which specifically aims at organising and/or financially supporting activities that stimulate network relations between the different types of actors that are involved in developing, implementing or embedding HLT technologies and encourage knowledge transfer between these actors. How do you judge the success of this part of the programme so far?

STEVIN has played an important role in encouraging network relations between the complete spectrum of players involved in developing, implementing and embedding HLT resources and technologies. This has been achieved in a number of highly visible ways:

1. By the nature of the calls for proposals, which have asked for joint projects involving industry and universities.
2. Through a number of major events to which representatives from across the spectrum of researchers, developers and users have been invited. Examples of such events are the brokerage events organized in connection with the calls for proposals, "Taal in Bedrijf"

('language@work'), STEVIN programme meetings, etc. These events were generally very successful and had high attendance numbers.

3. By supporting - both financially and in person - various networking events organized by research groups, companies and others. Various conferences⁴ (from the scientific oriented CLIN⁵ and InterSpeech2007⁶ towards more applied TaalinBedrijf⁷ and ICT-Delta⁸) were sponsored and/or used as a platform for networking and promotional activities. More details about the organized or sponsored networking events until 2007 (15 in total) can be found in the STEVIN Fact File.

Some eight publications on the STEVIN programme have appeared in several journals, in particular journals published by the programme bureau organisations and the journal of the NOTaS organization (DIXIT). See the section on question 6 and the STEVIN Fact File for more details.

The STEVIN programme also opened calls for proposals for companies to develop HLT-related demonstration applications (with or without universities). Companies could present interesting proposals for software applications where state-of-the-art language and speech technology were used in a new situation. If selected, the Nederlandse Taalunie would commission the consortium to develop the application. The application had to be realised within 15 months. Bureaucratic overhead was reduced to a minimum, the consortium would be paid a fair hourly wage for the number of hours estimated to build the application up to a maximum of € 100.000 for the completed demonstrator. It resulted in 14 demonstrator projects (see STEVIN Fact File). These applications cover a wide variety of LST for different groups in the Dutch and Flemish society (schools, counties, police, courtroom, children, etc.).

The major goals with these demonstration projects was to increase the visibility of language and speech technology by creating attractive demonstrators and show these demonstrators to governmental organizations, companies, educational organizations, researchers, the general public, etc., thus trying to increase demand to language and speech technology.

The results of these projects have been demonstrated at various events, including InterSpeech2007 in Antwerp, TaalinBedrijf, ICT-Delta, STEVIN programme meetings and at various other events.

From the beginning it was clear that a strong effort should be made in the dissemination of Language and Speech Technology for school children (age 17+) and bachelor students. The last years the number of language and speech technology-students has been decreased severely (for several reasons) and it becomes more and more difficult to find excellent (Dutch speaking) Master and PhD-students. Spreading the word for schoolchildren and bachelor students may help to increase the number of students in language and speech technology related areas and, as a side effect, make youngsters aware of something like LST. Since autumn-2007 the Program Committee solicited therefore for "educational" projects. One project was rewarded in 2007 and a new call for "education projects" has been open since 15-02-2008.

In combination with (at least some) of the projects granted in the third open call, we believe all the targets of the third goal have been addressed and have to a large extent also been reached in an excellent way.

We continue to think of new ways of achieving the stated goal, and one concrete suggestion to be implemented in 2008 is setting up master classes.

⁴ For a list of events supported by STEVIN networking grants see:

<http://taaluniversum.org/taal/technologie/stevin/netwerking/>

⁵ CLIN: <http://lands.let.ru.nl/-clin2007/sponsors.html>

⁶ INTERSPEECH2007: <http://www.interspeech2007.org/>

⁷ Taal in Bedrijf: <http://taaluniversum.org/taal/technologie/taalinbedrijf/>

⁸ ICT-Delta: <http://www.ictdelta.nu/>

Question 6

Are the STEVIN programme and its results visible to the (inter)national field of HLT research, the interested industry, and the interested public? Please provide some examples.

The visibility of the STEVIN programme and its (available) results can be considered significant, both at the national and international level. Also, research, industry and (interested) public are aware of the existence of STEVIN.

The STEVIN website generates a lot of general attention (www.stevin-tst.org). As of October 2006, when more accurate measuring of website visits started, the STEVIN site was visited on average 1,300 times per month. In 2007 the STEVIN site had 15,772 hits, i.e. approximately 45 per day. It is part of the Taalunieversum (www.taalunieversum.org), the website of the Dutch Language Union (Nederlandse Taalunie). The individual research project and demonstrator websites also provide information on STEVIN and the (expected) results. Furthermore, organisations like the Netherlands Organisation for Scientific Research (NWO; www.nwo.nl), NOTaS (www.notas.nl), the Dutch HLT Agency (TST-Centrale; www.tst.inl.nl) and SenterNovem (www.senternovem.nl) provide information about and links to the STEVIN website.

The answers to question 3 - (scientific) output - and 7 - (scientific) impact - provide information about and examples of the visibility of STEVIN within the (inter)national research community. Further proof of the visibility of STEVIN within the research community in the Netherlands and Flanders is the huge number of proposals (42 submitted proposals) submitted for the three open calls.

Given the number of demonstrator project proposals (40 submitted proposals), the industry is also very aware of the existence of STEVIN. Many Language and Speech Technology companies - mostly SMEs - are member of NOTaS, the Dutch Organisation for Language and Speech Technology. NOTaS and STEVIN often worked closely together in organising events (see also the answer to question 5) and in publishing STEVIN information. E.g. in DIXIT, NOTaS' illustrated magazine for members and their clients. One edition was a special STEVIN-themed edition (circulation: 4,250).

Other publications also paid attention to STEVIN, like articles about STEVIN in the "Innovatiekrant" (newspaper of the ministry of Economic Affairs, December 2007) and in "Hypothese" (an illustrated magazine published by NWO, November 2007). Various electronic newsletters paid attention to the STEVIN programme, news and (expected) results. E.g. the Language and Speech Technology newsletter published periodically by the Dutch Language Union (<http://taalunieversum.org/taal/technologie/nieuwsbrief>). A special STEVIN leaflet⁹ was created and handed out at events.

As shown in the answer to question 5, STEVIN organised and/or financially supported several events for research groups, companies and the (interested) public. In addition the HLT Agency (responsible for acquisition and IPR, management and maintenance, and distribution and service) has organised several events with presentations about STEVIN and expected STEVIN results.

It is difficult to provide metrics on the visibility of STEVIN at the level of the (interested) public. However, it is undeniably true that demonstrator projects have contributed to a higher visibility of STEVIN and Dutch language and speech technology in e.g. courts, government, large (commercial) organisations, police and municipalities (see STEVIN Fact File for more details on the demonstration projects). If the clients of the various language and speech technology organisations joined in NOTaS, the readers of the publications mentioned here and in the answers to other questions and visitors of the various websites with STEVIN information can all be considered (interested) public, then also their awareness of STEVIN must have been raised considerable.

At the time of writing, it is too early to evaluate the availability of STEVIN results other than the results of the finished demonstrator projects. The visibility of the (results of the) demonstrator

⁹ http://taalunieversum.org/taal/technologie/stevin/documenten/stevin_folder_staand.pdf

projects has been discussed above. Only few research projects have formally finished and therefore only few results have been made available for (re)use via the Dutch HLT Agency. However, some intermediate results have been made available to interested (research) parties. This happened in close cooperation between the project and the HLT Agency. Examples are the use of the D-COI text corpus by the University of Utrecht and the use of the Cornetto database by Princeton University U.S.A.

Question 7

What is the (scientific) impact of the STEVIN programme with respect to Dutch and Flemish HLT?

The scientific impact of a programme such as STEVIN can be measured in several different ways, but it will always be in terms of measurable differences between the situation at the start of the programme and the current situation. Such a comparison is hampered somewhat by the fact that the results of the so called *nulmeting* (the assessment of the situation at the start of the programme) have only been published in the middle of November 2007.

The *nulmeting* covers the following topics:

- Number of professors, faculty, support staff, PhD students and Master students
The latest numbers do not differ significantly from the numbers given in the report, but it is difficult to interpret the numbers. A not-so-positive interpretation is evidently that the impact of STEVIN has been difficult to attest. A more positive interpretation is that thanks to STEVIN the HLT field in the Low Countries has escaped the reductions in the number of staff members that we have seen in other subjects, especially in the Arts faculties. However, it has not yet led to an increase of structural (tenured) positions in HLT.
- Ratio of tenured staff, staff funded by (in)national research agencies and staff funded by contracts with industry
According to the report the total number of persons working in language & speech technology in the knowledge infrastructure was almost 170. Clearly, the budget available in STEVIN is not large enough to cause a visible difference in these numbers.
- Research topics
Not surprisingly, the research topics mentioned by the academic research groups at the start of the STEVIN programme overlap completely with the priorities formulated in the programme objectives. For this reason it is once again difficult to quantify the impact of STEVIN on the current situation. However, here too it is fair to say that funding from STEVIN has allowed the field to fend off the results of budget cuts in the Dutch universities.
- Resources
Three years after the start of the programme most of the resources aimed at in STEVIN are still under development; some have very recently been completed. Therefore, it is not reasonable to expect that new resources have already been used intensively for research and development. However, it should be pointed out that the Spoken Dutch Corpus, which was completed shortly after the start of STEVIN, has been used for a large number of research projects and that it is still being used intensively. Therefore, it is reasonable to assume that the resources under development in STEVIN will also be used intensively as soon as they become available to the R&D community. To substantiate this claim it can be mentioned that the SPRAAK software is playing a pivotal role in several projects and project proposals in STEVIN as well as in projects funded independent of STEVIN. Therefore, we are confident that STEVIN will make a real difference for future research in language & speech technology in the Dutch speaking countries.
- The last aspect addressed in the *nulmeting* report is the presence of the research labs in national and international networks. The report concludes that this representation is quite thinly spread and somewhat haphazard. The report does not mention collaborations between research groups that address related topics. It is, however, fair to say that already now STEVIN has strengthened

the collaboration (especially between groups in the Netherlands and corresponding groups in Flanders) considerably.

- So far, the impact of STEVIN on the number of conference and journal publications has been limited. It is even likely that high impact publication output (ACL, Computational Linguistics journal) has diminished. To a large degree this is due to the focus on resource creation rather than basic research or innovative applications; in addition, strategic research was largely focused on in the second call. Thus, by far the most important outlet for STEVIN papers is the bi-annual LREC conference.

Question 8

Is the organisational structure of the programme, namely the International Assessment Panel, HLT Board, programme committee, and programme bureau appropriate? Does the programme offer sufficient guarantees on transparent, impartial and objective proposal evaluation, project monitoring and decision taking in general? Are potential conflicts of interest adequately dealt with on all levels of the programme?

Organizational structure

The organisational structure of the programme was specifically designed to divide the programme management into complementary task packages that require specific and partly complementary expertise, to make different committees responsible for these tasks, and to setup a general procedure for the making of decisions.

The STEVIN Programme Committee (PC) is composed of Dutch and Flemish representatives of academia and industry together representing expertise in scientific research, application development and application introduction in the market. The PC is responsible for (1) developing the programme and formulating the calls for proposals that must realise the aims of the programme, (2) proposing a selection of awardable applications to the Board of the STEVIN programme, taking into account the IAP assessment, and (3) monitoring the progress made by the awarded projects.

The International Assessment Panel (IAP) is composed of eight internationally respected language and speech technologists and it is responsible for (1) providing an assessment of the scientific quality of the submitted proposals (including their anticipated contribution to the STEVIN programme objectives) to the PC, and (2) participating in the midterm assessment of the STEVIN programme.

The Board of the STEVIN programme (Board) consists of representatives of the sponsoring organisations, and a few scientific persons with experience in running large research and development programs. It is responsible for supervising the progress of the programme towards the programme goals and for ensuring the transparency and fairness of the assessment and ranking of project proposals by the PC. The Board is also responsible for supervising two other activities that are related to the STEVIN programme, i.e. the Dutch HLT Agency and the 'Makel en Schakel' activities carried out by the Nederlandse Taalunie. To guarantee a good communication between the Board and the PC, the chair person of the PC is also invited to the Board meetings.

The Programme Bureau (PB) consists of persons from NWO and SenterNovem, two organisations with a long experience in the practical implementation of large research projects. The PB also offers logistic assistance to the IAP, the PC and the Board, and in this way it acquires the survey that is needed to coordinate all the STEVIN activities. Obviously, the PB is also an excellent common source of information for the different committees.

Since the PC has always valued the project proposal assessments made by the IAP, since the Board has almost always followed the proposals made by the PC, and since the PB has always been able to carry out the recommendations made by the other committees, there have so far been no reasons for questioning the validity of the devised organisation. So far, there has been only one case in which the

Board seriously deviated from the PC recommendation that is worth mentioning. It will be discussed in more detail below.

Transparent, impartial and objective evaluation and monitoring of scientific projects

All calls for proposals have been widely announced via a variety of media (newsletters, NWO website, etc.). For the first and the second call, brokerage events have been organized to explain the nature of the calls and to give potential applicants the opportunity to meet with possible co-applicants. These brokerage events have been a great success, as can be seen from the number of participants in each brokerage event (160 and 165 participants respectively), and from the large number of proposals submitted in the three open STEVIN calls van R&D projects (52 full proposals, 68 if we include pre-proposals). In this way it was ensured that everyone with an interest in participating to the programme was aware of the opportunities STEVIN offers. This was the best guarantee for obtaining the best proposals and the best coverage of topics.

To make the evaluation impartial and objective, the assessment of the scientific quality of the project proposals was done by the IAP, a committee of independent and internationally highly respected speech and language technologists from all over Europe. The PC adhered to that assessment by not selecting any submitted project that was not highly ranked by the IAP. If the final PC ranking slightly differed from that of the IAP, this usually had to do with additional elements such as the responses of applicants to the assessment of the IAP, the balance between speech and language and between Dutch and Flemish involvement, the amount of collaboration implied, etc. The PC always justified the deviations of its ranking from that made by the IAP.

In analogy with the code of conduct used by the European Commission for its Framework Programme, a STEVIN Code of Conduct was formulated for the IAP and PC. All members were required to sign a declaration of conflict of interest and confidentiality and to formally indicate in which - if any - of the submitted projects they had any involvement. In doing so, the members committed themselves to strict confidentiality and impartiality concerning their tasks. If a member of the PC had a direct or indirect link with the project(s), or any other vested interest, or is in some way connected with the project(s), or has any other allegiance which impairs or threatens to impair his/her impartiality with respect to the project(s), the STEVIN Programme Bureau has ensured that this member did not participate in the review and ranking of the project(s) concerned.

To make the evaluation of project proposals transparent, the applicants were always informed about the assessments made by the IAP, and from the second call on, the applicants were given the opportunity to respond to these assessments. These responses were then taken into account by the PC in making the final ranking. Since any involvement of PC members in the ranked project proposals is unavoidable, the final decision on the funding of projects is made by the Board, which mainly consists of representatives of the Dutch and Flemish funding bodies and some HLT experts.

For the monitoring of projects, the PC asked the projects for half-yearly reports and for each project, two of its members (the so-called "Portefeuillehouders", abbreviated as PHs) are asked to perform a project progress control. A full check of the scientific quality and validity would be impossible to realize, because the PHs are often insufficiently knowledgeable in the very specialized topics addressed in the projects they have to monitor. Different projects are monitored by different PH teams, and the PHs are not directly related to the projects they have to monitor.

The monitoring of progress by the PC did not run optimally, because of several reasons. First, the exact nature of the task of the PHs was not clearly defined in the beginning, meaning that monitoring was done in different ways by different PHs. Second, the submissions of reports were not synchronized, but happened following a schedule that was related to the project start date. This made it difficult to centralise the monitoring of projects in particular PC meetings. So there was clearly room for some improvement in this respect.

The PC has therefore taken the following actions.

1. The PC has more clearly defined the task of the PHs, and developed a form to assist the PHs in their work.
2. The projects are now required to submit their half-yearly reports on fixed dates.

3. A PC meeting with project monitoring as an agenda item will be planned approximately a month after such a report submission date, so that the PHs have enough time to carry out their task in preparation of the PC meeting.
4. The monitoring of IPR related issues is left to the representative of the HLT Agency in the PC.

Another instrument that was set up for the monitoring of projects is the organisation of so-called STEVIN Programme Meetings. On such a workshop-like event, all projects present their activities, progress and results to representatives of the other running projects, to PC members, Board members and members of any other group within STEVIN programme organisation. So far, two of these meetings have been organized (in September 2006 and September 2007), and one is planned (September 2008).

Transparent, impartial and objective evaluation and monitoring of other projects

Apart from scientific projects, the STEVIN programme also funds demonstration projects and educational projects.

Demonstration projects are intended to stimulate the creation of new services and applications with existing speech and language technologies. Evaluation criteria focus on the potential contribution of a project to the visibility of services enabled by Language and Speech Technology. Another important evaluation criterion is the involvement of SMEs in these projects.

The proposals for demonstration projects have been evaluated and ranked by a small committee of mainly civil servants. The reasons for working with a civil servant committee were twofold: (1) avoiding conflict of interest; (2) the civil servant committee members have experience with these kinds of projects from other programmes and are especially experts in assessing the economic benefits of such projects.

The Programme Committee has had virtually no influence on the selection of application projects. There is only a so-called `sanity check` for the proposals that have been proposed for funding by the civil servant committee, but there usually was insufficient time to carry this out in a proper manner. Even though the reasons for working with a civil servant committee have been discussed with and accepted by the PC, at least some PC members have the opinion that the PC involvement (or, alternatively, involvement by the committee for accompanying actions) should have been more intense than it currently is.

With respect to the monitoring of demonstration projects, only few and incomplete reports on the results of these projects are as yet available. Improving the reporting of results is currently being worked on by the PC. In particular, the availability and location of demonstrators or running (demo) services is currently being inventoried.

Since `educating new experts' is one of the goals of the STEVIN accompanying measures, one call for **education projects** has been issued and others are planned for soliciting educational projects aiming at high school children (17+) and bachelor students. The formulation and evaluation of these calls are managed by the working group for accompanying actions, supported by a small committee of educational experts. As stated before, one project has been awarded.

Based on experiences with the evaluation of the first call for educational projects, some adaptations of the evaluation procedure are currently being made.

Deviation from the PC advice by the Board

As stated before, there has only been one serious difference of opinion between the PC and the Board that is worth mentioning. This originated from the fact that the advice of the PC concerning the Tender for creating the 500M word Written Dutch Corpus and the projects submitted in the third open call, was not followed by the Board. This difference of opinion had nothing to do with the scientific content of the advice, but it all had to do with the fact that at the time of the advice there was still no commitment for an amount of 628.000 Euro that had to be brought in by the Dutch sponsors to complete the funding of the STEVIN programme.

The advice of the PC stated that the proposal for creating the 500M word corpus of written Dutch had to be funded in first priority, and that the remaining money could then be spent to fund the first

ranked projects of the open call. However, at the time the Board had to make a decision about the granting of projects there was still no solution for the lacking funds, and the Board decided to use the already committed money to fund the 5 proposals from the third call and not the Tender for creating the 500M corpus. It did so for the following reasons. First, since the written corpus was the top priority of the BATAVO report and the STEVIN programme, claiming that there is no money for the creation of this priority would help the Board to raise the pressure on the Dutch authorities to find a solution for the lacking money. Second, the Board was unhappy with the under-representation of Flemish partners in the corpus consortium (and remained so despite extensive explanations of the reasons why Flemish research groups chose not to participate in an enterprise that did not offer room for research). The latter argument was especially pressing because it was the Dutch government that was responsible for the lacking funds. Although the decision taken by the Board was legitimate, it was highly unusual and deeply regretted by the PC, in particular due to the bad communication of this decision by the Board to the PC. A variety of actions was started up to obtain the lacking money. The most crucial of these was instigated by workers in the field, who wrote an open letter signed by more than 280 persons from the field to the Dutch responsible ministers, cc-ing the parliamentary members responsible for the Dutch language Union. This led to questions being asked in the Dutch parliament on this matter. Fortunately, as a result of these actions, the funds did become available in April 2008.

Question 9

Were the installed procedures for granting, launching and monitoring the projects and for transferring and approving the project results adequate?"

In R&D Call 1 applicants were asked to submit a full proposal directly. This was done for two reasons: (1) the amount of time from the opening of the call until the submission dead line was very short, since some of the money had to be committed before the end of 2004; (2) the budget for the call, and the size and duration of the projects in Call 1 were limited. The first call was certainly very successful (with 19 submitted and 5 awarded projects), and the granting procedure has been carried out successfully and in time. However, after the granting date, it still took much longer than originally aimed to have the projects actually start up. The reasons for this were manifold, but can be summarized by the statement that the whole constellation of the projects working under STEVIN was completely new to everybody involved, in various respects, and there was insufficient time to elaborate on these issues if the originally planned start dates were to be kept:

1. All awarded projects had multiple participants, as was intended since projects with multiple participants, preferably from both Flanders and the Netherlands, and both from knowledge institutions and from industry, were preferred. One participant had a special role as the project leader and project responsible, and as speaking partner to the programme bureau. This required making consortium agreements, with which neither the project participants nor the programme bureau had much experience.
2. The limited time available from the programme start and the granting of the first projects had left insufficient time for the programme bureau to prepare adequate granting letters, which contributed to delays in the projects start-up. The fact that a result obligation rather than a (at least in academia) more customary effort obligation was requested, also contributed to this.
3. The requirement that IPR issues be taken care of adequately was new to all involved parties and contributed to delays in the projects start-up. The project members were in a different relation to the programme coordinator (NTU) than was customary in Dutch and Flemish projects (results must be transferred to the NTU at the end of the project).
4. The IPR issues related to data from external partners (publishers etc.) also contributed to delays, at least for some projects.

These problems have been addressed in various ways. First, the consortium agreements of the first projects that had created one were turned into a consortium agreement template¹⁰ so that other projects, both from the first call and from later calls, could use this as a basis to start from. Second, a granting letter template¹¹ has been created, which could be used by other projects. This granting letter template has later been slightly adapted to accommodate new or unforeseen circumstances. Both the consortium agreement template and the granting letter template contain clauses guaranteeing the proper relation between NTU and the project participants. Finally, an IPR working group has been formed to assist with IPR problems in general, yielding a number of concrete results described earlier that facilitate dealing with these issues by the project participants, NTU, Dutch HLT Agency and external data providers.

The second call therefore faced less of these problems, also aided by the fact that more time was available from the call opening until the intended project start dates. In the second call, applicants were requested to first submit a pre-proposal, after which they were given an advice whether and how to proceed. All 34 project pre-proposals were commented upon by the PC. All project applicants were given the opportunity to respond to these comments, and all pre-proposals and the comments were discussed between the PC and representatives of the project pre-proposals. This was very much appreciated by the applicants, especially since the PC made concrete suggestions on how to improve the projects. Finally, 18 full proposals were submitted and again assessed and ranked by the IAP and the PC.

Already in the first call for proposals it became clear that the call as formulated by the Programme Committee invited excellent strategic research project proposals, but fewer proposals dedicated to creating new resources and tools or to application oriented research projects. To realise the development of two important resources, the PC therefore decided to launch, alongside the second open call for proposals, two specific calls for tenders, as described above.

When in the second open call, again virtually all top-ranking projects were of the second type (strategic research), the PC decided to amend the assessment criteria for the third and final open call in such a way that the submission of application-oriented proposals was stimulated. Furthermore, as it was felt that not enough proposals were of the highest quality, it was decided to not spend all the budget available for the second call, but reserve € 800.000 for a third call for tender for a large richly-annotated Dutch written corpus, which was the top priority in the BATAVO report and for which D-COI (awarded in the first call) was preparing the basis.

After extensive discussions in the Programme Committee and between the PC and the TST Board it was decided to use the unallocated budget for a third call for tender for a 500 M word corpus of written Dutch. It was first checked whether the results of D-COI were openly available to a sufficient degree, so that it was guaranteed that any consortium (not just the D-COI consortium) could make a potentially successful application. This check was carried out by external experts (with knowledge of the Dutch language, and not from the IAP or PC). The conclusion of this investigation was positive, so that the call could be launched.

For the three calls for tender four proposals were submitted in total. For the semantic lexicon some high quality competing proposals were made, among which it was very difficult both for the IAP and the PC to select one. For the other two tenders there was only one proposal submitted for each tender. For the SPRAAK project the consortium contained a good representation of the Dutch and Flemish groups working on speech recognition research. The composition of the SONAR project can be claimed to be representative as well, though more qualifications are in order here.

In preparing the third open call the PC asked two of its members (Daelemans and Martens) to prepare an assessment of the degree to which the priorities formulated at the start of the programme were covered by the funded projects. From the reports of Daelemans and Martens it became clear that most of the priorities for resources and strategic research were covered to the extent possible given

¹⁰ [consortium agreement template](http://taaluniversum.org/taal/technologie/stevin/documenten/model_CA.doc):

http://taaluniversum.org/taal/technologie/stevin/documenten/model_CA.doc

¹¹ [granting letter template](http://taaluniversum.org/taal/technologie/stevin/documenten/committingbrief_o&oproject.pdf) (in Dutch):

http://taaluniversum.org/taal/technologie/stevin/documenten/committingbrief_o&oproject.pdf

the limitations of the total budget. At the same time it was clear that there were substantial gaps in the area of application oriented research. For this reason application oriented research was explicitly emphasized in the third call (without excluding the possibility for submitting proposals for strategic research proposals). At the same time, the evaluation criteria for the PC and the IAP were extended to also justify the requirements of application oriented research proposals. In this call 15 projects proposals were submitted, and 5 of them awarded, all application-oriented research projects.

Summary

We can conclude that the installed procedures for granting, launching and monitoring the projects and for transferring and approving the project results are adequate. However these have had to large extent to be developed and refined during the programme. The major reason for this is that the whole set-up and constellation of the STEVIN programme is new, and, especially in the early phase, there was too little time to prepare procedures, forms etc. in advance. It is clear that this has led to delays in the start-up and finalization of certain projects, especially the first round projects. However, in the meantime most procedures have been elaborated and experiences gained have been incorporated, so that similar problems could be avoided for most of the projects from the other rounds.