

Audio Source Separation With a Single Sensor

Laurent Benaroya, Frédéric Bimbot, and Rémi Gribonval

Presented by Itai Ram

June 2007



Agenda

- Problem description
- Theoretical background
- Solution formulation
- Extension to Non-Gaussian densities
- Separation algorithms
- Training methods (optional)
- Performance evaluation
- References
- Supplementary (Personal extension)
- Performances comparison (TBD)
- Q&A



Problem description

Audio source separation – problem formulation:
Given observations:

$$x_i = \sum_{j=1}^m a_{i,j} s_j$$

Where:

x_i - the i 'th sensor

s_j - the j 'th source out of m sources

Estimate the sources: $\{\hat{s}_j\}_{j=1}^m$



Problem description (Cont')

- Two possible cases for this formulation:
 - $n \geq m$: The number of sensors is greater or equal to the number of sources. In this case the solution may be derived from the pseudo-inverse of the estimated matrix $A = (a_{i,j})$
 - $n < m$: The number of sensors is less than the number of sources . This case is also known as the *under-determined case*.
- In the article, the authors analyze an extreme case of the *under-determined case*:

$$x = s_1 + s_2$$



Theoretical background

- Maximum Likelihood estimation:

$$(\hat{s}_1, \hat{s}_2) = \arg \max_{s_1, s_2} p(x | s_1, s_2)$$

Where X is the observed data and S_1 and S_2 are the sources that need to be estimated.

- $p(x | s_1, s_2)$ is called the likelihood function.
- Since we would like to incorporate the sources' a priori knowledge, it is natural to consider in our case the MAP estimator by using the Bayesian formalism:

$$p(s_1, s_2 | x) \propto p(x | s_1, s_2) p(s_1) p(s_2)$$

While the two sources are assumed to be statistically independent and the information about them can be obtained as part of a prior training step [1]

[1] L. Benaroya, R. Gribonval, and F. Bimbot, "Non negative sparse representation for wiener based source separation with a single sensor," in Proc. ICASSP, Hong Kong, 2003, pp. 613–616.



Solution formulation

- Bayesian formulation of the Wiener filter
- Our observation:

$$x = s_1 + s_2 + n$$

Where:

$$s_1 \sim N(0, \Sigma_1), \quad s_2 \sim N(0, \Sigma_2) \quad \text{and} \quad n \sim N(0, \sigma^2 I)$$

- Relying on Bayes law we get the following MAP and MMSE estimators for the sources:

$$\hat{s}_1 = \Sigma_1 [\Sigma_1 + \Sigma_2 + \sigma^2 I]^{-1} x$$
$$\hat{s}_2 = \Sigma_2 [\Sigma_1 + \Sigma_2 + \sigma^2 I]^{-1} x$$



Solution formulation (Cont')

- In the case when S1 and S2 are stationary and circular processes (i.e., with a Toeplitz covariance matrix), it can be shown that the Wiener filtering can be interpreted as the following operator in the frequency domain:

$$\hat{s}_1[f] = \frac{\sigma_1^2[f]}{\sigma_1^2[f] + \sigma_2^2[f]} \cdot x[f]$$

$$\hat{s}_2[f] = \frac{\sigma_2^2[f]}{\sigma_1^2[f] + \sigma_2^2[f]} \cdot x[f]$$

- The *a priori* needed knowledge about the sources is in fact reduced to the covariance matrix – the PSD (Power Spectral Densities) in the frequency domain.



Extension to Non-Gaussian densities

- Audio signals are usually not Gaussian and are nonstationary (i.e., with frequency spectrum changing in time). Therefore, our approach needs to be expanded to other prior densities of the sources.
- In this paper, two main models are discussed – the Gaussian Mixture Model and the Gaussian Scaled Mixture Model. These models rely on the use of a dictionary of spectral shapes in order to deal with the non stationarity of the signals.
- In the supplementary part we will present additional statistical model for the sources.



The Gaussian Mixture Model

- Let's first define the Gaussian-centered (zero-mean) distribution:

$$g(y, \Sigma) = \frac{1}{(2\pi)^{N/2}} \frac{1}{\sqrt{\det \Sigma}} \exp\left[-\frac{1}{2} y^T \Sigma^{-1} y\right]$$

Where N refers to the dimension of the observation.

- Based on the following definition, we can now define the general structure of the *Gaussian Mixture Model*:

$$G\left(y, \{\varpi^{(i)}\}, \{\Sigma^{(i)}\}\right) = \sum_{i=1}^K \varpi^{(i)} g\left(y, \Sigma^{(i)}\right)$$

- This model consists of K Gaussians, each one with a probability in the mixture of $\varpi^{(i)}$ and:

$$\sum_{i=1}^K \varpi^{(i)} = 1$$



The Gaussian Mixture Model (Cont')

- For our objective of sources separation, the Gaussian Mixture Model allows us to deal with each one of the components, thus allowing us to use prior solution in order to calculate the a priori statistics of the sources.

$$p(s_1) = \sum_{i=1}^{K_1} \varpi_1^{(i)} \frac{\exp\left[-\frac{1}{2} s_1^T \Sigma_1^{(i)-1} s_1\right]}{(2\pi)^{N/2} \left|\det\left(\Sigma_1^{(i)}\right)\right|^{1/2}}$$

$$p(s_2) = \sum_{j=1}^{K_2} \varpi_2^{(j)} \frac{\exp\left[-\frac{1}{2} s_2^T \Sigma_2^{(j)-1} s_2\right]}{(2\pi)^{N/2} \left|\det\left(\Sigma_2^{(j)}\right)\right|^{1/2}}$$

With $\sum_{i=1}^{K_1} \varpi_1^{(i)} = 1$ and $\sum_{i=1}^{K_2} \varpi_2^{(i)} = 1$



The Gaussian Mixture Model (Cont')

- In order to use the previous formula, we will need to introduce hidden variables that are associated with the active component in the mixture.

$$p(s_i | q_i = k) = \frac{\exp\left[-\frac{1}{2} s_i^T \Sigma_1^{(k)-1} s_i\right]}{(2\pi)^{N/2} \left|\det\left(\Sigma_i^{(k)}\right)\right|^{1/2}}$$

$$p(q_i = k) = \varpi_i^{(k)}$$

- The estimators are therefore calculated conditionally to the hidden variables and there is a need to estimate them as well.
- Given the couple (q_1, q_2) we get the same estimators for the MMSE and the MAP:

$$E(s_1 | i, j, x) = \Sigma_1^{(i)} [\Sigma_1^{(i)} + \Sigma_2^{(j)} + \sigma^2 I]^{-1} x$$

$$E(s_2 | i, j, x) = \Sigma_2^{(j)} [\Sigma_1^{(i)} + \Sigma_2^{(j)} + \sigma^2 I]^{-1} x$$



The Gaussian Mixture Model (Cont')

- The next step is to estimate the active components. In order to do so, we will define the *a posteriori* density for the active components, given the observation:

$$\gamma_{i,j}(x) = P(i, j | x) \propto p(x | i, j) p(i) p(j)$$

- MAP: when the active components are not known, we can estimate them through the MAP estimator by calculating the couple (q_1, q_2) that maximizes $\gamma_{i,j}(x)$, yielding the following estimators:

$$E(s_1 | \hat{i}, \hat{j}, x) = \Sigma_1^{(\hat{i})} [\Sigma_1^{(\hat{i})} + \Sigma_2^{(\hat{j})} + \sigma^2 I]^{-1} x$$

$$E(s_2 | \hat{i}, \hat{j}, x) = \Sigma_2^{(\hat{j})} [\Sigma_1^{(\hat{i})} + \Sigma_2^{(\hat{j})} + \sigma^2 I]^{-1} x$$

- This approach is also referred to an adaptive Wiener filtering process.



The Gaussian Mixture Model (Cont')

- MMSE: another option to estimate the sources s_1 and s_2 , is by using the Bayes law with the a posteriori knowledge given through $\gamma_{i,j}(x)$

$$\begin{aligned} p(s_1 | x) &\propto \sum_{i,j} p(s_1 | x, i, j) p(q_1 = i, q_2 = j | x) \\ &\propto \sum_{i,j} p(s_1 | x, i, j) \gamma_{i,j}(x) \end{aligned}$$

- Finally the MMSE estimators for the sources are:

$$E(s_1 | x) = \sum_{i,j} \gamma_{i,j}(x) \cdot \Sigma_1^{(i)} [\Sigma_1^{(i)} + \Sigma_2^{(j)} + \sigma^2 I]^{-1} x$$

$$E(s_2 | x) = \sum_{i,j} \gamma_{i,j}(x) \cdot \Sigma_2^{(j)} [\Sigma_1^{(i)} + \Sigma_2^{(j)} + \sigma^2 I]^{-1} x$$



The Gaussian Scaled Mixture Model

- The major drawback of the GMM model is that it does not cover the case of different amplitudes for the same sound (i.e., a signal with the same frequencies content, but with different gain) which is essential when dealing with audio signal processing.
- If we would like to use the GMM model as described before, we will need as many Gaussian components as there are different possible amplitudes. This is quite limiting.
- The paper discussed presents an elaborated approach aims to solve this issue. This approach is referred to as the Gaussian Scaled Mixture Models.
- The concept of the GSMM is to mix scaled Gaussian densities according to a trained scaled parameters of the sound sources.



The Gaussian Scaled Mixture Model (Cont')

- We will define a Gaussian scaled density as follows:

$$g_a = \sqrt{a} \cdot g$$

- Where g is a Gaussian distributed random variable and a is a nonnegative scalar random variable.
- Therefore, given the scaled parameters of each component $\{a_i^1\}_{i=1}^{K_1}$, $\{a_j^2\}_{j=1}^{K_2}$ the GSMM takes the following form (sources' densities):

$$p(s_1 | a_1^1, \dots, a_{K_1}^1) = \sum_{i=1}^{K_1} \varpi_1^{(i)} \frac{\exp\left[-\frac{1}{2} s_1^T (a_i^1 \Sigma_1^{(i)})^{-1} s_1\right]}{(2\pi a_i^1)^{N/2} |\det(\Sigma_1^{(i)})|^{1/2}}$$

$$p(s_2 | a_1^2, \dots, a_{K_2}^2) = \sum_{j=1}^{K_2} \varpi_2^{(j)} \frac{\exp\left[-\frac{1}{2} s_2^T (a_j^2 \Sigma_2^{(j)})^{-1} s_2\right]}{(2\pi a_j^2)^{N/2} |\det(\Sigma_2^{(j)})|^{1/2}}$$



GSMM – Sources' estimators

- Assuming that the scaled factors $\{a_i^1\}_{i=1}^{K_1}, \{a_j^2\}_{j=1}^{K_2}$ are known and that the couple of the hidden variables of the active components ($q_1 = i, q_2 = j$) are known, the MAP and the MMSE estimators yield the same estimators:

$$E(s_1 | x, i, j, a_i^1, a_j^2) = a_i^1 \Sigma_1^{(i)} [a_i^1 \Sigma_1^{(i)} + a_j^2 \Sigma_2^{(j)} + \sigma^2 I]^{-1} x$$

$$E(s_2 | x, i, j, a_i^1, a_j^2) = a_j^2 \Sigma_2^{(j)} [a_i^1 \Sigma_1^{(i)} + a_j^2 \Sigma_2^{(j)} + \sigma^2 I]^{-1} x$$

- In order to develop the final MMSE estimator we will now define, similar to the GMM case, the *a posteriori* density for the active components, given the observation and the scaled factors:

$$\gamma_{i,j|a_i^1, a_j^2}(x) = P(i, j | x, a_i^1, a_j^2) \propto p(x | i, j, a_i^1, a_j^2) p(i) p(j)$$



GSMM – MMSE estimator

- In order to calculate the MMSE estimator, we will need to integrate over all active components and scaled factors, yielding the following estimator:

$$E(s_1 | x) = \sum_{i=1}^{K_1} \sum_{j=1}^{K_2} \int_{a_i^1} \int_{a_j^2} \gamma_{i,j,a_i^1,a_j^2}(x) \times E(s_1 | x, i, j, a_i^1, a_j^2) p_0(a_i^1) p_0(a_j^2) da_i^1 da_j^2$$

- To avoid integrating the above expression over all scale factors, a more simple solution is to estimate them by using the ML approach and then setting the values estimated separately.

$$\left(\hat{a}_i^1, \hat{a}_j^2 \right) = \arg \max_{a_1 \geq 0, a_2 \geq 0} \gamma_{i,j,a_i^1,a_j^2}(x)$$

- This approach yields the following estimator:

$$E(s_1 | x) = \sum_{i=1}^{K_1} \sum_{j=1}^{K_2} \gamma_{i,j,\hat{a}_i^1,\hat{a}_j^2}(x) \cdot \hat{a}_i^1 \sum_{i=1}^i [\hat{a}_i^1 \sum_{i=1}^i + \hat{a}_j^2 \sum_{j=1}^j + \sigma^2 I]^{-1} x$$



Separation algorithms

- Our objective is to separate two independent sources that are locally stationary. Therefore, it is natural to work in the STFT domain. In the STFT domain we get the following expression:

$$S_x(t, f) = S_{s_1}(t, f) + S_{s_2}(t, f) + S_b(t, f)$$

With covariance matrices of the sources - $\Sigma_1^{(i)}, \Sigma_2^{(j)}$ that are assumed to be diagonal.



Separation algorithm – GMM model

- Characterization of each of the sources – s1 and s2, by a GMM model, i.e., finding:

$$\{\varpi_1^{(i)}, \sigma_1^{(i)}(f)\}_{1 \leq i \leq K_1} \quad \text{and} \quad \{\varpi_2^{(j)}, \sigma_2^{(j)}(f)\}_{1 \leq j \leq K_2}$$

- For each frame index t , calculate the *a posteriori* conditional density of the active Gaussians components:

$$\gamma_{i,j}(t) \propto \varpi_1^{(i)} \varpi_2^{(j)} \prod_f g(|Sx(t, f)|, \sigma_1^{(i)}(f)^2 + \sigma_2^{(j)}(f)^2 + \sigma^2)$$

- MAP estimator:

$$\hat{i}, \hat{j} = \arg \max_{i,j} \gamma_{i,j}(t)$$

$$\hat{S}_{s_1}(t, f) = \frac{\sigma_1^{(\hat{i})}(f)^2}{\sigma_1^{(\hat{i})}(f)^2 + \sigma_2^{(\hat{j})}(f)^2 + \sigma^2} Sx(t, f)$$

$$\hat{S}_{s_2}(t, f) = \frac{\sigma_2^{(\hat{j})}(f)^2}{\sigma_1^{(\hat{i})}(f)^2 + \sigma_2^{(\hat{j})}(f)^2 + \sigma^2} Sx(t, f)$$

- MMSE estimator:

$$\hat{S}_{s_1}(t, f) = \sum_{i=1}^{K_1} \sum_{j=1}^{K_2} \gamma_{i,j}(t) \frac{\sigma_1^{(i)}(f)^2}{\sigma_1^{(i)}(f)^2 + \sigma_2^{(j)}(f)^2 + \sigma^2} Sx(t, f)$$

$$\hat{S}_{s_2}(t, f) = \sum_{i=1}^{K_1} \sum_{j=1}^{K_2} \gamma_{i,j}(t) \frac{\sigma_2^{(j)}(f)^2}{\sigma_1^{(i)}(f)^2 + \sigma_2^{(j)}(f)^2 + \sigma^2} Sx(t, f)$$



Separation algorithm – GSMM model

- The separation algorithm for the GSMM model is similar to that of the GMM model except for an additional estimation of the gain factors:
- Characterization of each of the sources – s1 and s2, by a GMM model.
- For each frame index time t :

- Estimate gain factors using ML method to achieve $\{\hat{a}_i^1(t), \hat{a}_j^2(t)\}$

- Calculate:
$$\gamma_{i,j}(t) \propto \bar{w}_1^{(i)} \bar{w}_2^{(j)} \prod_f g(|Sx(t, f)|, a_1^{(i)} \sigma_1^{(i)}(f)^2 + a_2^{(j)} \sigma_2^{(j)}(f)^2 + \sigma^2)$$

- MAP estimator:
$$\hat{i}, \hat{j} = \arg \max_{i,j} \gamma_{i,j}(t)$$

$$\hat{S}_{s_1}(t, f) = \frac{a_1^{(\hat{i})} \sigma_1^{(\hat{i})}(f)^2}{a_1^{(\hat{i})} \sigma_1^{(\hat{i})}(f)^2 + a_2^{(\hat{j})} \sigma_2^{(\hat{j})}(f)^2 + \sigma^2} Sx(t, f)$$

$$\hat{S}_{s_2}(t, f) = \frac{a_2^{(\hat{j})} \sigma_2^{(\hat{j})}(f)^2}{a_1^{(\hat{i})} \sigma_1^{(\hat{i})}(f)^2 + a_2^{(\hat{j})} \sigma_2^{(\hat{j})}(f)^2 + \sigma^2} Sx(t, f)$$

- MMSE estimator:

$$\hat{S}_{s_1}(t, f) = \sum_{i=1}^{K_1} \sum_{j=1}^{K_2} \gamma_{i,j}(t) \frac{a_1^{(i)} \sigma_1^{(i)}(f)^2}{a_1^{(i)} \sigma_1^{(i)}(f)^2 + a_2^{(j)} \sigma_2^{(j)}(f)^2 + \sigma^2} Sx(t, f)$$

$$\hat{S}_{s_2}(t, f) = \sum_{i=1}^{K_1} \sum_{j=1}^{K_2} \gamma_{i,j}(t) \frac{a_2^{(j)} \sigma_2^{(j)}(f)^2}{a_1^{(i)} \sigma_1^{(i)}(f)^2 + a_2^{(j)} \sigma_2^{(j)}(f)^2 + \sigma^2} Sx(t, f)$$



Performance evaluation

- In order to analyze the separation algorithms performances, the authors refer to the estimated sources as an orthogonal projection of the estimators on the real sources:

$$\hat{s}_1 = \alpha_1 s_1 + \alpha_2 s_2 + n_1$$

$$\hat{s}_2 = \beta_1 s_1 + \beta_2 s_2 + n_2$$

- They define the following criteria:

$$SIR_1 = 20 \log_{10} \left| \frac{\alpha_1}{\alpha_2} \frac{\|s_1\|}{\|s_2\|} \right| \quad SAR_1 = 20 \log_{10} \frac{\|\hat{s}_1 - n_1\|}{\|n_1\|}$$

$$SIR_2 = 20 \log_{10} \left| \frac{\beta_2}{\beta_1} \frac{\|s_2\|}{\|s_1\|} \right| \quad SAR_2 = 20 \log_{10} \frac{\|\hat{s}_2 - n_2\|}{\|n_2\|}$$

- The SIR (Source to Interference Ratio) is a way to measure the residual of the other source in the estimation of each source, whereas the SAR (Source to Artifacts Ratio) is an estimate of the amount of distortion in each estimated signal.



Performance evaluation (Cont')

- The experiment discussed in the article consists of 45sec of training the two sources – a piano and drums which are assumed to be uncorrelated (-16dB correlation). The signals are then mixed and are estimated during a 15sec time frame.
- Both the GMM and the GSMM models are evaluated using various numbers of components in their mixtures.

SIR

Criterion		MMSE		MAP	
State	Source	GMM	GSMM	GMM	GSMM
Wiener	Piano	8.7	4.7	15.5	4.4
Wiener	Drums	6.7	19.6	0.7	13.8
4	Piano	10.5	11.0	20.0	10.5
4	Drums	9.7	18.5	2.8	18.2
8	Piano	11.0	11.1	20.0	10.7
8	Drums	11.3	18.2	4.0	18.1
16	Piano	11.8	12.9	16.9	12.6
16	Drums	11.9	16.9	4.2	16.3

SAR

Criterion		MMSE		MAP	
State	Source	GMM	GSMM	GMM	GSMM
Wiener	Piano	7.8	12.3	1.8	12.4
Wiener	Drums	5.8	0.0	14.7	-0.5
4	Piano	8.4	8.7	2.0	8.6
4	Drums	5.9	4.0	9.6	3.7
8	Piano	8.9	8.5	3.4	8.7
8	Drums	5.9	4.0	8.9	4.0
16	Piano	8.1	8.6	3.5	8.4
16	Drums	5.4	5.3	8.3	5.0



Training methods (optional)

- The training of the sources is necessary in order to retrieve the *a priori* knowledge about the sources needed for the separation algorithm (Building the *codebook* for the sources).
- The authors present [1] some training algorithms that can be used in our case:
 - Randomized algorithm
 - Correlation based algorithm
 - Additive mixture based algorithm
- Two more iterative algorithms:
 - *K*-means
 - E-M (Expectation – Maximization)



References

- L. Benaroya, R. Blouet, C. Févotte and I. Cohen, “SINGLE SENSOR SOURCE SEPARATION BASED ON WIENER FILTERING AND MULTIPLE WINDOW STFT”
- Ephraim, Y., Malah, D., 1984. “Speech enhancement using a minimum mean-square error short-time spectral amplitude estimator,” IEEE Trans. Acoust. Speech Signal Process. ASSP-32 (6), 1109–1121.
- L. Benaroya, R. Gribonval, and F. Bimbot, “Non negative sparse representation for wiener based source separation with a single sensor,” in ICASSP, Hong Kong, 2003, pp. 613–616.



Supplementary

- In this part we present a different statistical model for the sources and their new estimators.
- The development of this solution is based on Prof. Israel Cohen paper on “*Speech enhancement using super-Gaussian speech models and noncausal a priori SNR estimation,*” *Speech Communication*, Vol. 47, No. 3, Nov. 2005, pp. 336-350.



The Laplacian model

- We will propose another statistical model for the problem, according to which one of the sources – the speech - has a Laplacian prior density and the other source – the instrument – remains with Gaussian prior density.
- An MMSE estimator in the STFT domain for both sources will be developed and presented.
- Future work should be done to compare the performances of the separation algorithm between the GMM/GSMM model and the Laplacian model.



The Laplacian model (Cont')

- Let's consider the problem of speech enhancement, in which we have a noisy signal such that:

$$y = x + d$$

While x denotes the speech and d the uncorrelated additive noise signal. Both processes are centered and independent.

- Applying the STFT on the following observation we get:

$$Y(k, l) = X(k, l) + D(k, l)$$



The Laplacian model (Cont')

- The MMSE estimator in the STFT domain is obtained by:

$$\hat{X} = E(X | Y)$$

which can also be written as:

$$\hat{X} = G(\xi, \gamma)Y$$

With:

$$\xi(k, l) \triangleq \frac{\lambda_x(k, l)}{\lambda_D(k, l)} \qquad \gamma(k, l) \triangleq \frac{Y^2(k, l)}{\lambda_D(k, l)}$$

$$\lambda_x(k, l) \triangleq E\{|X(k, l)|^2\}, \quad \lambda_D(k, l) \triangleq E\{|D(k, l)|^2\}$$

G is called the gain function.



Problem formulation

- Let's consider the following observation:

$$x = s_1 + s_2$$

Where now:

$$s_1 \sim L(0, \lambda_L), \quad s_2 \sim N(0, \lambda_G)$$

- The Laplacian density:

$$p(X) = \frac{1}{\sqrt{\lambda_L}} \exp\left(-\frac{2|X|}{\lambda_L}\right)$$

- In this particular case, we will assume that s_1 is the speech and s_2 is the musical instrument.
- Our objective: Given x , derive an MMSE estimators for both sources (i.e, the previously described gain functions).



MMSE estimator for the Laplacian source

- Based on results from speech enhancement, we get the following MMSE estimator for the Laplacian source – s_1 :

$$\hat{s}_1 = E(s_1 | x) = \frac{2}{L_{\rho+} - L_{\rho-}} \times \frac{L_{\rho+} \operatorname{erfcx}(L_{\rho+}) - L_{\rho-} \operatorname{erfcx}(L_{\rho-})}{\operatorname{erfcx}(L_{\rho+}) + \operatorname{erfcx}(L_{\rho-})}$$

Where:

$$L_{\rho\pm} = \sqrt{\frac{\lambda_G}{\lambda_L}} \pm \sqrt{\frac{X^2}{\lambda_G}}$$

And the scaled complementary error function is defined by:

$$\operatorname{erfcx}(x) \triangleq e^{x^2} \frac{2}{\sqrt{\pi}} \int_x^{\infty} e^{-t^2} dt$$



MMSE estimator for the Gaussian source

- Based on the problem formulation and after calculating \hat{s}_1 , the MMSE estimator for s_2 is simply calculated by:

$$\hat{s}_2 = x - E[s_1 | x] = x - \hat{s}_1$$

- Where \hat{s}_2 is the estimation for the second source, with the Gaussian prior density (assumed to be the musical instrument).



What is left?

- To complete the work, the following need to be done:
 - Achieve simulations results with the new model (SIR/SAR).
 - Compare simulations results with those of the GMM/GSMM models.
 - Results evaluation.
 - Define training methods.
 - TBD



Q & A

