

CONVERGENCE OF DISTRIBUTED ASYNCHRONOUS LEARNING VECTOR QUANTIZATION ALGORITHMS.

Benoît PATRA ^{a,b}

^a LSTA

Université Pierre et Marie Curie – Paris VI

Tour 15-25, 4 place Jussieu
75252 Paris cedex 05, France

benoit.patra@upmc.fr

^b LOKAD SAS

10 rue Philippe de Champaigne
75013 Paris, France

benoit.patra@lokad.com

Abstract

Motivated by the problem of effectively executing clustering algorithms on very large data sets, we address a model for large scale distributed clustering methods. To this end, we briefly recall some standards on the quantization problem and some results on the almost sure convergence of the Competitive Learning Vector Quantization (CLVQ) procedure. A general model for linear distributed asynchronous algorithms well adapted to several parallel computing architectures is also discussed. Our approach brings together this scalable model and the CLVQ algorithm, and we call the resulting technique the Distributed Asynchronous Learning Vector Quantization algorithm (DALVQ). An in-depth analysis of the almost sure convergence of the DALVQ algorithm is performed. A striking result is that we prove that the multiple versions of the quantizers distributed among the processors in the parallel architecture asymptotically reach a consensus almost surely. Furthermore, we also show that these versions converge almost surely towards the same nearly optimal value for the quantization criterion.

Keywords — k -means, vector quantization, distributed, asynchronous, stochastic optimization, scalability, distributed consensus.

1 Introduction

Distributed algorithms arise in a wide range of applications, including telecommunications, distributed information processing, scientific computing, real time process control and many others. Parallelization is one of the most promising ways to harness greater computing resources, whereas building faster serial computers is increasingly expensive and also faces some physical limits such as transmission speeds and miniaturization. One of the challenges proposed for Machine Learning is to build scalable applications that quickly process large amounts of data in sophisticated ways. Building such large scale algorithms attacks several problems in a distributed framework, such as communication delays in the network or numerous problems caused by the lack of shared memory.

Clustering algorithms are one of the primary tools of unsupervised learning. From a practical perspective, clustering plays an outstanding role in data mining applications such as text mining, web analysis, marketing, medical diagnostics, computational biology and many others. Clustering is a separation of data into groups of similar objects. As clustering represents the data with fewer clusters, there is a necessary loss of certain fine details, but simplification is achieved. The popular Competitive Learning Vector Quantization (CLVQ) algorithm (see Gersho and Gray [18]) provides a technique for building reliable clusters characterized by their prototypes. As pointed out by Bottou in [11], the CLVQ algorithm can also be viewed as the on-line version of the widespread Lloyd's method (see Lloyd's [25] for the definition) which is referred to as batch k -means in [11]. The CLVQ also belongs to the class of stochastic gradient descent algorithms (for more information on stochastic gradient descent procedures we refer the reader to Benveniste et al. [5]).

The analysis of parallel stochastic gradient procedures in a Machine Learning context has recently received a great deal of attention (see for instance Langford et al. [38] and Mac Donald et al. [27]). In the present paper, we go further by introducing a model that brings together the original CLVQ algorithm and the comprehensive theory of asynchronous parallel linear algorithms developed by Tsitsiklis [37], Tsitsiklis et al. [36] and Bertsekas and Tsitsiklis [6]. The resulting model will be called Distributed Asynchronous Learning Vector Quantization (DALVQ for short). At a high level, the DALVQ algorithm parallelizes several executions of the CLVQ method concurrently on different processors while the results of these algorithms are broadcast through the distributed framework asynchronously and efficiently.

Here, the term processor refers to any computing instance in a distributed architecture (see Bullo et al. [13, Chapter 1] for more details).

For a given time span, our parallel DALVQ algorithm is able to process much more data than a single processor execution of the CLVQ procedure. Moreover, DALVQ is also asynchronous. This means that local algorithms do not have to wait at preset points for messages to become available. This allows some processors to compute faster and execute more iterations than others, and it also allows communication delays to be substantial and unpredictable. The communication channels are also allowed to deliver messages out of order, that is, in a different order than the one in which they were transmitted. Asynchronism can provide two major advantages. First, a reduction of the synchronization penalty, which could bring a speed advantage over a synchronous execution. Second, for potential industrialization, asynchronism has greater implementation flexibility. Tolerance to system failures and uncertainty can also be increased. As in the case with any on-line algorithm, DALVQ also deals with variable data loads over time. In fact, on-line algorithms avoid tremendous and non scalable batch requests on all data sets. Moreover, with an on-line algorithm, new data may enter the system and be taken into account while the algorithm is already running.

As a striking result, we prove that multiple versions of the quantizers, distributed among the processors in a parallel architecture, asymptotically reach a consensus almost surely. Furthermore, we also show that these versions converge almost surely towards (the same) nearly optimal value for the quantization criterion. These convergence results are similar in spirit to the most satisfactory almost sure convergence theorem for the CLVQ algorithm obtained by Pagès in [29].

The paper is organized as follows. In Section 2 we review some standard facts on the clustering problem. We extract the relevant material from Pagès [29] without proof, thus making our exposition self-contained. In Section 3 we give a brief exposition of the mathematical framework for parallel asynchronous gradient methods introduced by Tsitsiklis et al. in [36] and Bertsekas and Tsitsiklis [36, 6]. The results of Blondel et al. [8] on the asymptotic consensus in asynchronous parallel averaging problems are also recalled. In Section 4, our main results are stated and proved.

2 Quantization and CLVQ algorithm

2.1 Overview

Let μ be a probability measure on \mathbb{R}^d with finite second-order moment. The quantization problem consists in finding a “good approximation” of μ by a set of κ vectors of \mathbb{R}^d called quantizers (or prototypes). Throughout the document the κ quantizers will be seen as the components of a $(\mathbb{R}^d)^\kappa$ -dimensional vector $w = (w_1, \dots, w_\kappa)$. To measure the correctness of a quantization scheme given by w , one introduces a cost function called distortion, defined by

$$C_\mu(w) = \frac{1}{2} \int_{\mathbb{R}^d} \min_{1 \leq \ell \leq \kappa} \|\mathbf{z} - w_\ell\|^2 d\mu(\mathbf{z}).$$

Under some minimal assumptions, the existence of an optimal quantizer vector $w^\circ \in \operatorname{argmin}_{w \in (\mathbb{R}^d)^\kappa} C_\mu(w)$ has been established by Pollard in [31] (see also Sabin and Gray [34, Appendix 2]).

In a statistical context, the distribution μ is only known through n independent random observations $\mathbf{z}_1, \dots, \mathbf{z}_n$ drawn according to μ . Denote by μ_n the empirical distribution based on $\mathbf{z}_1, \dots, \mathbf{z}_n$, that is, for every Borel subset A of \mathbb{R}^d

$$\mu_n(A) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\mathbf{z}_i \in A\}}.$$

Much attention has been devoted to the convergence study of the quantization scheme provided by the empirical minimizers

$$w_n^\circ \in \operatorname{argmin}_{w \in (\mathbb{R}^d)^\kappa} C_{\mu_n}(w).$$

The almost sure convergence of $C_\mu(w_n^\circ)$ towards $\min_{w \in (\mathbb{R}^d)^\kappa} C_\mu(w)$ was proved by Pollard in [30, 31] and Abaya and Wise in [2]. Rates of convergence and nonasymptotic performance bounds have been considered by Pollard [32], Chou [14], Linder et al. [24], Bartlett et al. [4], Linder [23, 35], Antos [1] and Antos et al. [3]. Convergence results have been established by Biau et al. in [7] where μ is a measure on a Hilbert space. It turns out that the minimization of the empirical distortion is a computationally hard problem. As shown by Inaba et al. in [20], the computational complexity of this minimization problem is exponential in the number of quantizers κ and the dimension of the data d . Therefore, exact computations are untractable for most of the practical applications.

Based on this, our goal in this document is to investigate effective methods that produce accurate quantizations with data samples. One of the most popular procedure is Lloyd’s algorithm (see Lloyd [25]) sometimes referred to as batch k -means. A convergence theorem for this algorithm is provided by Sabin and Gray in [34]. Another celebrated quantization algorithm is the Competitive Learning Vector Quantization (CLVQ), also called on-line k -means. The latter acronym outlines the fact that data arrive over time while the execution of the algorithm and their characteristics are unknown until their arrival times. The main difference between the CLVQ and the Lloyd’s algorithm is that the latter run in batch training mode. This means that the whole training set is presented before performing an update, whereas the CLVQ algorithm uses each item of the training sequence at each update.

The CLVQ procedure can be seen as a stochastic gradient descent algorithm. In the more general context of gradient descent methods, one cannot hope for the convergence of the procedure towards global minimizers with a non convex objective function (see for instance Benveniste et al. [5]). In our quantization context, the distortion mapping C_μ is not convex (see for instance Graf and Luschgy [19]). Thus, just as in Lloyd’s method, the iterations provided by the CLVQ algorithm converge towards local minima of C_μ .

Assuming that the distribution μ has a compact support and a bounded density with respect to the Lebesgue measure, Pagès states in [29] a result regarding the almost sure consistency of the CLVQ algorithm towards critical points of the distortion C_μ . The author shows that the set of critical points necessarily contains the global and local optimal quantizers. The main difficulties in the proof arise from the fact that the gradient of the distortion is singular on κ -tuples having equal components and the distortion function C_μ is not coercive. This explains why standard theories for stochastic gradient algorithm do not apply in this context.

2.2 The quantization problem, basic properties

In the sequel, we denote by \mathcal{G} the closed convex hull of $\text{supp}(\mu)$, where $\text{supp}(\mu)$ stands for the support of the distribution. Observe that, with this notation, the distortion mapping is the function $C : (\mathbb{R}^d)^\kappa \rightarrow [0, \infty)$ defined by

$$C(w) \triangleq \frac{1}{2} \int_{\mathcal{G}} \min_{1 \leq \ell \leq \kappa} \|\mathbf{z} - w_\ell\|^2 d\mu(\mathbf{z}), \quad w = (w_1, \dots, w_\kappa) \in (\mathbb{R}^d)^\kappa. \quad (2.1)$$

Throughout the document, with a slight abuse of notation, $\|\cdot\|$ means both the Euclidean norm of \mathbb{R}^d or $(\mathbb{R}^d)^\kappa$. In addition, the notation \mathcal{D}_*^κ stands for the set of all vector of $(\mathbb{R}^d)^\kappa$ with pairwise distinct components, that is,

$$\mathcal{D}_*^\kappa \triangleq \{w \in (\mathbb{R}^d)^\kappa \mid w_\ell \neq w_k \text{ if and only if } \ell \neq k\}.$$

Under some extra assumptions on μ , the distortion function can be rewritten using space partition set called Voronoï tessellation.

Definition 2.1 *Let $w \in (\mathbb{R}^d)^\kappa$, the Voronoï tessellation of \mathcal{G} related to w is the family of open sets $\{W_\ell(w)\}_{1 \leq \ell \leq \kappa}$ defined as follows:*

- If $w \in \mathcal{D}_*^\kappa$, for all $1 \leq \ell \leq \kappa$,

$$W_\ell(w) = \left\{ v \in \mathcal{G} \mid \|w_\ell - v\| < \min_{k \neq \ell} \|w_k - v\| \right\}.$$

- If $w \in (\mathbb{R}^d)^\kappa \setminus \mathcal{D}_*^\kappa$, for all $1 \leq \ell \leq \kappa$,

- if $\ell = \min \{k \mid w_k = w_\ell\}$,

$$W_\ell(w) = \left\{ v \in \mathcal{G} \mid \|w_\ell - v\| < \min_{w_k \neq w_\ell} \|w_k - v\| \right\}$$

- otherwise, $W_\ell(w) = \emptyset$.

As an illustration, Figure 1 shows Voronoï tessellations associated to a vector $w \in ([0, 1] \times [0, 1])^{50}$ whose components have been drawn independently and uniformly. This figure also highlights a remarkable property of the cell borders, which are portions of hyperplanes (see Graf and Luschgy [19]).

Observe that if $\mu(H)$ is zero for any hyperplane H of \mathbb{R}^d (a property which is sometimes referred to as strong continuity) then, using Definition 2.1, it is easy to see that the distortion takes the form:

$$C(w) = \frac{1}{2} \sum_{\ell=1}^{\kappa} \int_{W_\ell(w)} \|z - w_\ell\|^2 d\mu(z), \quad w \in (\mathbb{R}^d)^\kappa. \quad (2.2)$$

The following assumption will be needed throughout the paper. This assumption is similar to the peak power constraint (see Chou [14] and Linder [35]). Note that most of the results of this subsection are still valid if μ satisfies the weaker strong continuity property.




Figure 1: Voronoi tessellation of 50 points of \mathbb{R}^2 drawn uniformly in a square.

Assumption 2.1 (Compact Supported Density) *The probability measure μ has a bounded density with respect to the Lebesgue measure on \mathbb{R}^d . Moreover, the support of μ is equal to its compact convex hull \mathcal{G} .*

The next proposition states the differentiability of the distortion C , and provides an explicit formula for the gradient ∇C whenever the distortion is differentiable.

Proposition 2.1 (Pagès [29]) *Under Assumption 2.1, the distortion C is continuously differentiable at every $w = (w_1, \dots, w_\kappa) \in \mathcal{D}_*^\kappa$. Furthermore, for all $1 \leq \ell \leq \kappa$,*

$$\nabla_\ell C(w) = \int_{W_\ell(w)} (w_\ell - \mathbf{z}) d\mu(\mathbf{z}).$$

Some necessary conditions on the location of the minimizers of C can be derived from its differentiability properties. Therefore, Proposition 2.2 below states that the minimizers of C have parted components and that they are contained in the support of the density. Thus, the gradient is well defined and these minimizers are necessarily some zeroes of ∇C . For the sequel it is convenient to let $\overset{\circ}{A}$ be the interior of any subset A of $(\mathbb{R}^d)^\kappa$.

Proposition 2.2 (Pagès [29]) *Under Assumption 2.1, we have*

$$\operatorname{argmin}_{w \in (\mathbb{R}^d)^\kappa} C(w) \subset \operatorname{argminloc}_{w \in \mathcal{G}^\kappa} C(w) \subset \overset{\circ}{\mathcal{G}}^\kappa \cap \{\nabla C = 0\} \cap \mathcal{D}_*^\kappa,$$

where $\operatorname{argminloc}_{w \in \mathcal{G}^\kappa} C(w)$ stands for the set of local minimizers of C over \mathcal{G}^κ .

For any $\mathbf{z} \in \mathbb{R}^d$ and $w \in (\mathbb{R}^d)^\kappa$, define

$$H(\mathbf{z}, w) \triangleq \left((w_\ell - \mathbf{z}) \mathbb{1}_{\{\mathbf{z} \in W_\ell(w)\}} \right)_{1 \leq \ell \leq \kappa}. \quad (2.3)$$

On \mathcal{D}_*^κ , the function H may be interpreted as an observation of the gradient. With this notation, Proposition 2.1 states that

$$\nabla C(w) = \int_{\mathcal{G}} H(\mathbf{z}, w) d\mu(\mathbf{z}), \quad w \in \mathcal{D}_*^\kappa. \quad (2.4)$$

Let $\complement A$ stands for the complementary in $(\mathbb{R}^d)^\kappa$ of a subset $A \subset (\mathbb{R}^d)^\kappa$. Clearly, for all $w \in \complement \mathcal{D}_*^\kappa$, the mapping $H(\cdot, w)$ is integrable. Therefore, ∇C can be extended on $(\mathbb{R}^d)^\kappa$ via the formula

$$h(w) \triangleq \int_{\mathcal{G}} H(\mathbf{z}, w) d\mu(\mathbf{z}), \quad w \in (\mathbb{R}^d)^\kappa. \quad (2.5)$$

Note however that the function h , which is sometimes called the average function of the algorithm, is not continuous.

Remark 2.1 *Under Assumption 2.1, a computation for all $w \in \mathcal{D}_*^\kappa$ of the Hessian matrix $\nabla^2 C(w)$ can be deduced from Theorem 4 of Fort and Pagès [16]. In fact, the formula established in this theorem is valid for cost functions which are more complex than C (they are associated to Kohonen Self Organizing Maps, see Kohonen [21] for more details). In Theorem 4, letting $\sigma(k) = \mathbb{1}_{\{k=0\}}$, provides the result for our distortion C . The resulting formula shows that h is singular on $\complement \mathcal{D}_*^\kappa$ and, consequently, that this function cannot be Lipschitz on \mathcal{G}^κ .*

2.3 Convergence of the CLVQ algorithm

The problem of finding a reliable clustering scheme for a dataset is equivalent to find optimal (or at least nearly optimal) minimizers for the mapping C . A minimization procedure by a usual gradient descent method cannot be implemented as long as ∇C is unknown. Thus, the gradient is approximated

by a single example extracted from the data. This leads to the following stochastic gradient descent procedure

$$w(t+1) = w(t) - \varepsilon_{t+1} H(\mathbf{z}_{t+1}, w(t)), \quad t \geq 0, \quad (2.6)$$

where $w(0) \in \overset{\circ}{\mathcal{G}}^\kappa \cap \mathcal{D}_*^\kappa$ and $\mathbf{z}_1, \mathbf{z}_2 \dots$ are independent observations distributed according to the probability measure μ .

The algorithm defined by the iterations (2.6) is known as the CLVQ algorithm in the data analysis community. It is also called the Kohonen Self Organizing Map algorithm with 0 neighbor (see for instance Kohonen [21]) or the on-line k -means procedure (see MacQueen [26] and Bottou [10]) in various fields related to statistics. As outlined by Pagès in [29], this algorithm belongs to the class of stochastic gradient descent methods. However, the almost sure convergence of this type of algorithm cannot be obtained by general tools such as Robbins-Monro method (see Robbins and Monro [33]) or the Kushner-Clark's Theorem (see Kushner and Clark [22]). Indeed, the main difficulty essentially arises from the non convexity of the function C , its non coercivity and the singularity of h at $\overset{\circ}{\mathcal{D}}_*^\kappa$ (we refer the reader to [29, Section 6] for more details).

The following assumption set is standard in a gradient descent context. It basically upraises constraints on the decreasing speed of the sequence of steps $\{\varepsilon_t\}_{t=0}^\infty$.

Assumption 2.2 (Decreasing steps) *The $(0, 1)$ -valued sequence $\{\varepsilon_t\}_{t=0}^\infty$ satisfies the following two constraints:*

1. $\sum_{t=0}^\infty \varepsilon_t = \infty$.
2. $\sum_{t=0}^\infty \varepsilon_t^2 < \infty$.

An examination of identities (2.6) and (2.3) reveals that if $\mathbf{z}_{t+1} \in W_{\ell_0}(w(t))$, where $\ell_0 \in \{1, \dots, M\}$ then,

$$w_{\ell_0}(t+1) = (1 - \varepsilon_{t+1}) w_{\ell_0}(t) + \varepsilon_{t+1} \mathbf{z}_{t+1}. \quad (2.7)$$

The component $w_{\ell_0}(t+1)$ can be viewed as the image of $w_{\ell_0}(t)$ by a \mathbf{z}_{t+1} -centered homothety with ratio $1 - \varepsilon_{t+1}$ (Figure 2 provides an illustration of this fact). Thus, under Assumptions 2.1 and 2.2, the trajectories of $\{w(t)\}_{t=0}^\infty$ stay in $\overset{\circ}{\mathcal{G}}^\kappa \cap \mathcal{D}_*^\kappa$. More precisely, if

$$w(0) \in \overset{\circ}{\mathcal{G}}^\kappa \cap \mathcal{D}_*^\kappa$$

then,

$$w(t) \in \overset{\circ}{\mathcal{G}}^\kappa \cap \mathcal{D}_*^\kappa, \quad t \geq 0, \text{ almost surely (a.s.)}$$

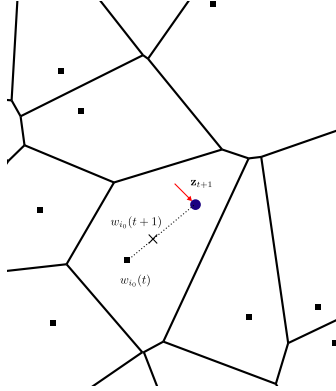


Figure 2: Drawing of a portion of a 2-dimensional Voronoi tessellation. For $t \geq 0$, if $\mathbf{z}_{t+1} \in W_{\ell_0}(w(t))$ then $w_\ell(t+1) = w_\ell(t)$ for all $\ell \neq \ell_0$ and $w_{\ell_0}(t+1)$ lies in the segment $[w_{\ell_0}(t), \mathbf{z}_{t+1}]$. The update of the vector $w_{\ell_0}(t)$ can also be viewed as a \mathbf{z}_{t+1} -centered homothety with ratio $1 - \varepsilon_{t+1}$.

Although ∇C is not continuous some regularity can be obtained. To this end, we need to introduce the following materials. For any $\delta > 0$ and any compact set $L \subset \mathbb{R}^d$, let the compact set $L_\delta^\kappa \subset (\mathbb{R}^d)^\kappa$ be defined as

$$L_\delta^\kappa \triangleq \left\{ w \in L^\kappa \mid \min_{k \neq \ell} \|w_\ell - w_k\| \geq \delta \right\}. \quad (2.8)$$

The next lemma that states on the regularity of ∇C will prove to be extremely useful in the proof of Theorem 2.2 and throughout Section 4.

Lemma 2.1 *Assume that μ satisfies Assumption 2.1 and let L be a compact set of \mathbb{R}^d . Then, there is some constant P_δ such that for all w and v in L_δ^κ with $[w, v] \subset \mathcal{D}_*^\kappa$,*

$$\|\nabla C(w) - \nabla C(v)\| \leq P_\delta \|w - v\|.$$

The following lemma, called G-Lemma in [29] is an easy-to-apply convergence results on stochastic algorithms. It is particularly adapted to the present situation where the average function of the algorithm h is singular.

Theorem 2.1 (G-Lemma, Fort and Pagès [17]) *Assume that the following conditions are satisfied:*

1. $\sum_{t=1}^{\infty} \varepsilon_t = \infty$ and $\varepsilon_t \xrightarrow[t \rightarrow \infty]{} 0$.
2. The sequences $\{w(t)\}_{t=0}^{\infty}$ and $\{h(w(t))\}_{t=0}^{\infty}$ are bounded a.s.
3. The series $\sum_{t=0}^{\infty} \varepsilon_{t+1} (H(\mathbf{z}_{t+1}, w(t)) - h(w(t)))$ converge a.s. in $(\mathbb{R}^d)^\kappa$.
4. There exists a lower semi-continuous function $G : (\mathbb{R}^d)^\kappa \rightarrow [0, \infty)$ such that

$$\sum_{t=0}^{\infty} \varepsilon_{t+1} G(w(t)) < \infty, \quad a.s.$$

Then, there exists a connected component Ξ of $\{G = 0\}$ such that

$$\text{dist}(w(t), \Xi) \xrightarrow[t \rightarrow \infty]{} 0, \quad a.s.$$

The interest of the G-Lemma depends upon the choice of G . In our context, a suitable lower semi-continuous function is \widehat{G} defined by

$$\widehat{G}(w) \triangleq \liminf_{v \in \mathcal{G}^\kappa \cap \mathcal{D}_*^\kappa, v \rightarrow w} \|\nabla C(v)\|^2, \quad w \in \mathcal{G}^\kappa. \quad (2.9)$$

The next theorem is, as far as we know, the first almost sure convergence theorem for the stochastic algorithm CLVQ.

Theorem 2.2 (Pagès [29]) *Under Assumptions 2.1 and 2.2, on the event*

$$\left\{ \liminf_{t \rightarrow \infty} \text{dist}(w(t), \mathcal{C}\mathcal{D}_*^\kappa) > 0 \right\}, \quad \text{one has}$$

$$\text{dist}(w(t), \Xi_\infty) \xrightarrow[t \rightarrow \infty]{} 0, \quad a.s.,$$

where Ξ_∞ is some connected component of $\{\nabla C = 0\}$.

The proof is an application of the above G-Lemma with the mapping \widehat{G} defined by equation (2.9). Theorem 2.2 states that the iterations of the CLVQ necessarily converge towards some critical points (zeroes of ∇C). From Proposition 2.2 we deduce that the set of critical points necessarily contains optimal quantizers. Remind that without more assumption than $w(0) \in \overset{\circ}{\mathcal{G}}^\kappa \cap \mathcal{D}_*^\kappa$, we have already discussed the fact that the components of $w(t)$ are almost surely parted for all $t \geq 0$. Thus, it is easily seen that the two following events are equal

$$\left\{ \liminf_{t \rightarrow \infty} \text{dist}(w(t), \mathcal{C}\mathcal{D}_*^\kappa) > 0 \right\} = \left\{ \inf_{t \geq 0} \text{dist}(w(t), \mathcal{C}\mathcal{D}_*^\kappa) > 0 \right\}.$$

Some results are provided by Pagès in [29] for asymptotically stuck components but, as pointed out by the author, they are less satisfactory.

3 General distributed asynchronous algorithm

3.1 Model description

Let $s(t)$ be any $(\mathbb{R}^d)^\kappa$ -valued vector and consider the following iterations on a vector $w \in (\mathbb{R}^d)^\kappa$

$$w(t+1) = w(t) + s(t), \quad t \geq 0. \quad (3.1)$$

Here, the iterations described by identity (3.1) can only be performed by a single computing entity. Therefore, if the computations of the vectors $s(t)$ are relatively time consuming then, not many iterations can be achieved for a given time span. Consequently, a parallelization of this computing scheme should be investigated. The aim of this section is to discuss a precise mathematical description of a distributed asynchronous model for the iterations (3.1). This model for distributed computing was originally proposed by Tsitsiklis et al. in [36] and was revisited in Bertsekas and Tsitsiklis [6, Section 7.7].

Assume that we dispose of a distributed architecture with M computing entities called processors (or agents, see for instance Bullo et al. [13]). Each processor is labeled, for simplicity of notation, by a natural number $i \in \{1, \dots, M\}$. Throughout the paper, we will add the superscript i on the variables possessed by the processor i . In the model we have in mind, each processor has a buffer where its current version of the iterated vector is kept, i.e a local memory. Thus, for agent i such iterations are represented by the $(\mathbb{R}^d)^\kappa$ -valued sequence $\{w^i(t)\}_{t=0}^\infty$.

Let $t \geq 0$ denote the current time. For any pair of processors $(i, j) \in \{1, \dots, M\}^2$, the value kept by agent j and available for agent i at time t is not necessarily the most recent one, $w^j(t)$, but more probably and outdated one, $w^j(\tau^{i,j}(t))$, where the deterministic time instant $\tau^{i,j}(t)$ satisfy $0 \leq \tau^{i,j}(t) \leq t$. Thus, the quantity $t - \tau^{i,j}(t)$ represents a communication and possibly other types of delay.

We insist on the fact that there is a distinction between “global” and “local” time. The time variable we refer above to as t corresponds to a global clock. Such a global clock is needed only for analysis purposes. The processors work without knowledge of this global clock. They have access to a local clock or to no clock at all.

The algorithm is initialized at $t = 0$, where each processor $i \in \{1, \dots, M\}$ has an initial version $w^i(0) \in (\mathbb{R}^d)^\kappa$ in its buffer. We define the general distributed asynchronous algorithm by the following iterations

$$w^i(t+1) = \sum_{j=1}^M a^{i,j}(t) w^j(\tau^{i,j}(t)) + s^i(t), \quad i \in \{1, \dots, M\} \text{ and } t \geq 0. \quad (3.2)$$

The model can be interpreted as follows: at time $t \geq 0$, processor i receives messages from other processors containing $w^j(\tau^{i,j}(t))$. Processor i incorporates these new vectors by forming a convex combination and incorporates the vector $s^i(t)$ resulting from its own “local” computations. The coefficients $a^{i,j}(t)$ are nonnegative numbers which satisfy the constraint

$$\sum_{j=1}^M a^{i,j}(t) = 1, \quad i \in \{1, \dots, M\} \text{ and } t \geq 0. \quad (3.3)$$

As the combining coefficients $a^{i,j}(t)$ depend on t , the network communication topology is sometimes referred to as time-varying. The sequences $\{\tau^{i,j}(t)\}_{t=0}^\infty$ need not to be known in advance by any processor. In fact, their knowledge is not required to execute iterations defined by equation (3.2). Thus, we do not necessarily dispose of a shared global clock or synchronized local clocks at the processors. The difference $t - \tau^{i,j}(t)$ can be seen as a communication delay. This is a modeling of some aspects of the network: latency and bandwidth finiteness.

As for now the descent terms $\{s^i(t)\}_{t=0}^\infty$ will be arbitrary $(\mathbb{R}^d)^\kappa$ -valued sequences. In Section 4, when we define the Distributed Asynchronous Learning Vector Quantization (DALVQ), the definition of the descent terms will be made more explicit.

3.2 The agreement algorithm

This subsection is devoted to a short survey of the results, found by Blondel et al. in [8], for a natural simplification of the general distributed asynchronous algorithm (3.2). This simplification is called agreement algorithm by Blondel et al. and is defined by

$$x^i(t+1) = \sum_{j=1}^M a^{i,j}(t) x^j(\tau^{i,j}(t)), \quad i \in \{1, \dots, M\} \text{ and } t \geq 0. \quad (3.4)$$




Figure 3: Illustration of the time delays introduced in the general distributed asynchronous algorithm. Here, there are $M = 4$ different processors with their own computations of the vectors $w^{(i)}$, $i \in \{1, 2, 3, 4\}$. Three arbitrary values of the global time t are represented (t_1 , t_2 and t_3), with $\tau^{i,i}(t_k) = t_k$ for all $i \in \{1, 2, 3, 4\}$ and $1 \leq k \leq 3$. The dashed arrows head towards the versions available at time t_k for an agent $i \in \{1, 2, 3, 4\}$ represented by the tail of the arrow.

where $x^i(0) \in (\mathbb{R}^d)^k$. An observation of these equations reveals that they are similar to iterations (3.2), the only difference being that all descent terms equal 0.

In order to analyze the convergence of the agreement algorithm (3.4), Blondel et al in [8] define two sets of assumptions that enforce some weak properties on the communication delays and the network topology. As shown in [8], if the assumptions contained in one of these two set hold, then the distributed versions of the agreement algorithm, namely the x^i , reach an asymptotical consensus. This latter statement means that there exists a vector x^* (independent of i) such that

$$x^i(t) \xrightarrow[t \rightarrow \infty]{} x^*, \quad i \in \{1, \dots, M\}.$$

The agreement algorithm (3.4) is essentially driven by the communication

times $\tau^{i,j}(t)$ assumed to be deterministic but do not need to be known *a priori* by the processors. The following Assumption 3.1 essentially ensures, in its third statement, that the communication delays $t - \tau^{i,j}(t)$ are bounded. This assumption prevents some processor from taking into account some arbitrarily old values computed by others processors. Assumption 3.1 1. is just a convention: when $a^{i,j}(t) = 0$ the value $\tau^{i,j}(t)$ has no effect on the update. Assumption 3.1 2. is rather natural because processors have access to their own most recent value.

Assumption 3.1 (Bounded communication delays) 1. If $a^{i,j}(t) = 0$ then, $\tau^{i,j}(t) = t$, $(i, j) \in \{1, \dots, M\}^2$ and $t \geq 0$,

2. $\tau^{i,i}(t) = t$, $i \in \{1, \dots, M\}$ and $t \geq 0$.

3. There exists a positive integer B_1 such that

$$t - B_1 < \tau^{i,j}(t) \leq t, \quad (i, j) \in \{1, \dots, M\}^2 \text{ and } t \geq 0.$$

The next Assumption 3.2 states that the value possessed by agent i at time $t + 1$, namely $x^i(t + 1)$, is a weighted average of its own value and the values that it has just received from other agents.

Assumption 3.2 (Convex combination and threshold) There exists a positive constant $\alpha > 0$ such that the following three properties hold:

1. $a^{i,i}(t) \geq \alpha$, $i \in \{1, \dots, M\}$ and $t \geq 0$.

2. $a^{i,j}(t) \in \{0\} \cup [\alpha, 1]$, $(i, j) \in \{1, \dots, M\}^2$ and $t \geq 0$.

3. $\sum_{j=1}^M a^{i,j}(t) = 1$, $i \in \{1, \dots, M\}$ and $t \geq 0$.

Let us mention one particular relevant case for the choice of the combining coefficients $a^{i,j}(t)$. Let $i \in \{1, \dots, M\}$ and $t \geq 0$, the set

$$N^i(t) \triangleq \{j \in \{1, \dots, M\} \in \{1, \dots, M\} \mid a^{i,j}(t) \neq 0\}$$

corresponds to the set of agents whose version is taken into account by processor i at time t . For all $(i, j) \in \{1, \dots, M\}^2$ and $t \geq 0$, the weights $a^{i,j}(t)$ are defined by

$$a^{i,j}(t) = \begin{cases} 1/\#N^i(t) & \text{if } j \in N^i(t); \\ 0 & \text{otherwise;} \end{cases}$$

where $\#A$ denotes the cardinal of any finite set A . The above definition on the combining coefficients appears to be relevant for practical implementations of the model DALVQ introduced in Section 4. For a discussion on

others special interest cases regarding the choices of the coefficients $a^{i,j}(t)$ we refer the reader to [8].

The communication patterns, sometimes referred to as the network communication topology, can be expressed in terms of oriented graph. For a thorough introduction to graph theory, see Jungnickel [15].

Definition 3.1 (Communication graph) *Let us fix $t \geq 0$, the communication graph at time t , $(\mathcal{V}, E(t))$, is defined by*

- *the set of vertices \mathcal{V} is formed by the set of processors $\mathcal{V} = \{1, \dots, M\}$,*
- *the set of edges $E(t)$ is defined via the relationship*

$$(j, i) \in E(t) \text{ if and only if } a^{i,j}(t) > 0.$$

Assumption 3.3 is a minimal condition required for a consensus among the processors. More precisely, it states that for any pair of agents $(i, j) \in \{1, \dots, M\}^2$ there is a sequence of communications where the values computed by agent i will influence (directly or indirectly) the future values kept by agent j .

Assumption 3.3 (Graph connectivity) *The graph $(\mathcal{V}, \cup_{s \geq t} E(s))$ is strongly connected for all $t \geq 0$.*

Finally, we define two supplementary assumptions. The combination of one of the two following assumptions with the three previous ones will ensure the convergence of the agreement algorithm. As mentioned above, if Assumption 3.3 holds then, there is a communication path between any pair of agents. Assumption 3.4 below expresses the fact that there is a finite upper bound for the length of such paths.

Assumption 3.4 (Bounded communication intervals) *If i communicates with j an infinite number of times then, there is a positive integer B_2 such that*

$$(i, j) \in E(t) \cup E(t+1) \cup \dots \cup E(t+B_2-1), \quad t \geq 0.$$

Assumption 3.5 is a symmetry condition: if agent $i \in \{1, \dots, M\}$ communicates with agent $j \in \{1, \dots, M\}$ then, j has communicated or will communicate with i during the time interval $(t - B_3, t + B_3)$ where $B_3 > 0$.

Assumption 3.5 (Symmetry) *There exists some $B_3 > 0$ such that whenever $(i, j) \in E(t)$, there exists some τ that satisfies $|t - \tau| < B_3$ and $(j, i) \in E(\tau)$.*

To shorten a little bit the notation, we set

$$(\mathbf{AsY})_1 \equiv \begin{cases} \text{Assumption 3.1;} \\ \text{Assumption 3.2;} \\ \text{Assumption 3.3;} \\ \text{Assumption 3.4.} \end{cases} \quad (\mathbf{AsY})_2 \equiv \begin{cases} \text{Assumption 3.1;} \\ \text{Assumption 3.2;} \\ \text{Assumption 3.3;} \\ \text{Assumption 3.5;} \end{cases}$$

We are now in a position to state the main result of this section. The Theorem 3.1 expresses the fact that, for the agreement algorithm, a consensus is asymptotically reached by the agents.

Theorem 3.1 (Blondel et al. [8]) *Under the set of Assumptions $(\mathbf{AsY})_1$ or $(\mathbf{AsY})_2$, there is a consensus vector $x^* \in (\mathbb{R}^d)^k$ (independent of i) such that*

$$\lim_{t \rightarrow \infty} \|x^i(t) - x^*\| = 0, \quad i \in \{1, \dots, M\}.$$

Besides, there exist $\rho \in [0, 1)$ and $L > 0$ such that

$$\|x^i(t) - x^i(\tau)\| \leq L\rho^{t-\tau}, \quad i \in \{1, \dots, M\} \text{ and } t \geq \tau \geq 0.$$

3.3 Asymptotic consensus

This subsection is devoted to the analysis of the general distributed asynchronous algorithm (3.2). For this purpose, the study of the agreement algorithm defined by equations (3.4) will be extremely fruitful. The following lemma states that the version possessed by agent $i \in \{1, \dots, M\}$ at time $t \geq 0$, namely $w^i(t)$, depends linearly on the others initialization vectors $w^j(0)$ and the descent subsequences $\{s^j(\tau)\}_{\tau=-1}^{t-1}$, where $j \in \{1, \dots, M\}$.

Lemma 3.1 (Tsitsiklis [37]) *For all $(i, j) \in \{1, \dots, M\}^2$ and $t \geq 0$, there exists a real-valued sequence $\{\phi^{i,j}(t, \tau)\}_{\tau=-1}^{t-1}$ such that*

$$w^i(t) = \sum_{j=1}^M \phi^{i,j}(t, -1) w^j(0) + \sum_{\tau=0}^{t-1} \sum_{j=1}^M \phi^{i,j}(t, \tau) s^j(\tau).$$

For all $(i, j) \in \{1, \dots, M\}^2$ and $t \geq 0$, the real-valued sequences $\{\phi^{i,j}(t, \tau)\}_{\tau=-1}^{t-1}$ do not depend on the value taken by the descent terms $s^i(t)$. The real numbers $\phi^{i,j}(t, \tau)$ are determined by the sequences $\{\tau^{i,j}(\tau)\}_{\tau=0}^t$ and $\{a^{i,j}(\tau)\}_{\tau=0}^t$. These last sequences are unknown in general, but some useful qualitative properties can be derived, as expressed in Lemma 3.2 below.

Lemma 3.2 (Tsitsiklis [37]) *For all $(i, j) \in \{1, \dots, M\}^2$, let $\{\phi^{i,j}(t, \tau)\}_{\tau=-1}^{t-1}$ be the sequences defined in Lemma 3.1.*

1. Under Assumption 3.2,

$$0 \leq \phi^{i,j}(t, \tau) \leq 1, \quad (i, j) \in \{1, \dots, M\}^2 \text{ and } t > \tau \geq -1.$$

2. Under Assumptions $(\mathbf{AsY})_1$ or $(\mathbf{AsY})_2$, we have:

(a) For all $(i, j) \in \{1, \dots, M\}^2$ and $\tau \geq -1$, the limit of $\phi^{i,j}(t, \tau)$ as t tends to infinity exists and is independent of j . It will be denoted $\phi^i(\tau)$.

(b) There exists some $\eta > 0$ such that

$$\phi^i(\tau) > \eta, \quad i \in \{1, \dots, M\} \text{ and } \tau \geq -1.$$

(c) There exist a constant $A > 0$ and $\rho \in (0, 1)$ such that

$$|\phi^{i,j}(t, \tau) - \phi^i(\tau)| \leq A\rho^{t-\tau}, \quad (i, j) \in \{1, \dots, M\}^2 \text{ and } t > \tau \geq -1.$$

Take $t' \geq 0$ and assume that the agents stop performing update after time t' , but keep communicating and merging the results. This means that $s^j(t) = 0$ for all $t \geq t'$. Then equations (3.2) write

$$w^i(t+1) = \sum_{j=1}^M a^{i,j}(t) w^j(\tau^{i,j}(t)), \quad i \in \{1, \dots, M\} \text{ and } t \geq t'.$$

If Assumptions $(\mathbf{AsY})_1$ or $(\mathbf{AsY})_2$ are satisfied then, Theorem 3.1 shows that there is a consensus vector, depending on the time instant t' . Let us denote this vector by $w^*(t')$ (see Figure 4). However, Lemma 3.2 provides a better way to define the sequence $\{w^*(t)\}_{t=0}^\infty$ as shown in Definition 3.4. Note that this definition does not involve any assumption on the descent terms.

Definition 3.2 (Agreement vector) Assume that Assumptions $(\mathbf{AsY})_1$ or $(\mathbf{AsY})_2$ are satisfied. The agreement vector sequence $\{w^*(t)\}_{t=0}^\infty$ is defined by

$$w^*(t) \triangleq \sum_{j=1}^M \phi^j(-1) w^j(0) + \sum_{\tau=0}^{t-1} \sum_{j=1}^M \phi^j(\tau) s^j(\tau), \quad t \geq 0.$$

It is noteworthy that the agreement vector sequence w^* satisfies the following recursion formula

$$w^*(t+1) = w^*(t) + \sum_{j=1}^M \phi^j(t) s^j(t), \quad t \geq 0. \quad (3.5)$$

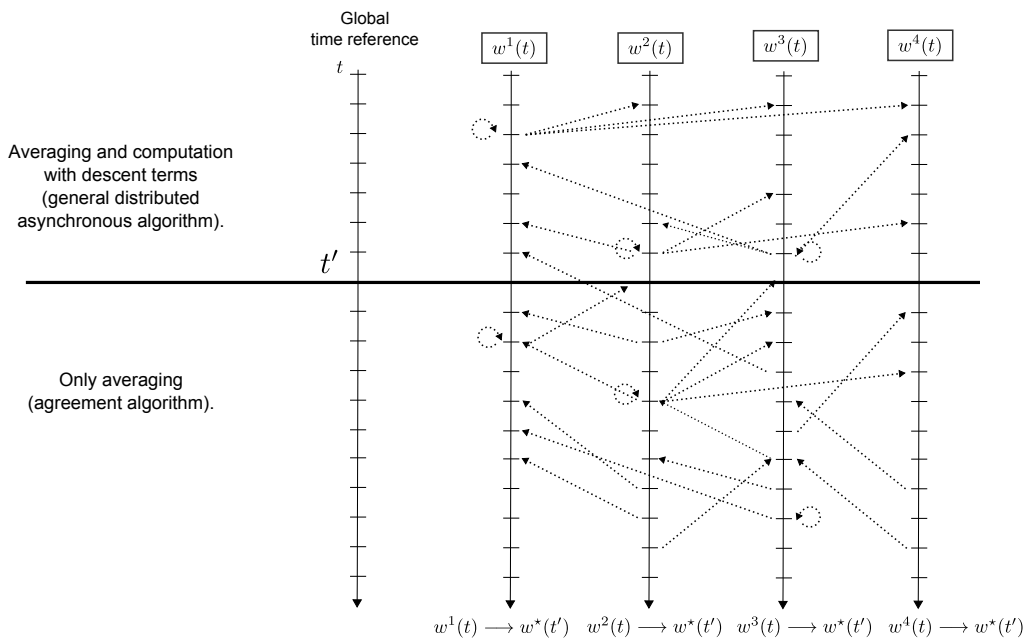


Figure 4: The agreement vector at time t' , $w^*(t')$ corresponds to the common value asymptotically achieved by all processors if computations integrating descent terms have stopped after t' , i.e, $s^j(t) = 0$ for all $t \geq t'$.

4 Distributed asynchronous learning vector quantization

4.1 Introduction, model presentation

From now on, and until the end of the paper, we assume that one of the two set of assumptions $(\mathbf{AsY})_1$ or $(\mathbf{AsY})_2$ holds, as well as the compact-supported density Assumption 2.1. In addition, we will also assume that $0 \in \mathcal{G}$. For the sake of clarity, all the proofs of the main theorems as well as the lemmas needed for these proofs have been postponed at the end of the paper, in Annex.

Tsitsiklis in [37], Tsitsiklis et al in [36] and Bertsekas and Tsitsiklis in [6] studied distributed asynchronous stochastic gradient optimization algorithms. In this series of publications, for the distributed minimization of a cost function $F : (\mathbb{R}^d)^K \rightarrow \mathbb{R}$, the authors considered the general distributed asynchronous algorithm defined by equation (3.2) with specific choices for stochastic descent terms s^i . Using the notation of Section 3, the algorithm writes

$$w^i(t+1) = \sum_{j=1}^M a^{i,j}(t) w^j(\tau^{i,j}(t)) + s^i(t), \quad i \in \{1, \dots, M\} \text{ and } t \geq 0,$$

with stochastic descent terms $s^i(t)$ satisfying

$$\mathbb{E} \{ s^i(t) \mid s^j(\tau), j \in \{1, \dots, M\} \text{ and } t > \tau \geq 0 \} = -\varepsilon_{t+1}^i \nabla F(w^i(t)), \\ i \in \{1, \dots, M\} \text{ and } t \geq 0. \quad (4.1)$$

where $\{\varepsilon_t^i\}_{t=0}^\infty$ are decreasing steps sequences. The definition of the descent terms in [6, 36] is more general than the one appearing in equation (4.1). We refer the reader to Assumption 3.2 and 3.3 in [36] and Assumption 8.2 in [6] for the precise definition of the descent terms in [6, 36]. As discussed in Section 2, the CLVQ algorithm is also a stochastic gradient descent procedure. Unfortunately, the results from Tsitsiklis et al. in [36, 37, 6] do not apply with our distortion function, C , since the authors assume that F is continuously differentiable and ∇F is Lipschitz. Therefore, the aim of this section is to extend the results of Tsitsiklis et al. to the context of vector quantization and on-line clustering.

We first introduce the Distributed Asynchronous Learning Vector Quantization (DALVQ) algorithm. To prove its almost sure consistency, we will need

an Asynchronous G-Lemma, which is inspired from the G-Lemma, Theorem 2.1, presented in Section 2. This theorem may be seen as an easy-to-apply tool for the almost sure consistency of a distributed asynchronous system where the average function is not necessary regular. Our approach sheds also some new light on the convergence of distributed asynchronous stochastic gradient descent algorithms. Precisely, Proposition 8.1 in [36] claims that $\liminf_{t \rightarrow \infty} \|\nabla F(w^i(t))\| = 0$ while our main Theorem 4.2 below states that $\lim_{t \rightarrow \infty} \|\nabla C(w^i(t))\| = 0$. However, there is a price to pay for this more precise result with the non Lipschitz gradient ∇C . Similarly to Pagès [29], who assumes that the trajectory of the CLVQ algorithm has almost surely asymptotically parted components (see Theorem 2.2 in Section 2), we will suppose that the agreement vector sequence has, almost surely, asymptotically parted component trajectories.

Recall that the goal of the DALVQ is to provide a well designed distributed algorithm that processes quickly (in term of wall clock time) very large data sets to produce accurate quantization. The data sets (or streams of data) are distributed among several queues sending data to the different processors of our distributed framework. Thus, in this context the sequence $\mathbf{z}_1^i, \mathbf{z}_2^i, \dots$ stands for the data available for processor, where $i \in \{1, \dots, M\}$. The random variables

$$\mathbf{z}_1^1, \mathbf{z}_2^1, \dots, \mathbf{z}_1^2, \mathbf{z}_2^2, \dots$$

are assumed to be independent and identically distributed according to μ .

In the definition of the CLVQ procedure (2.6), the term $H(\mathbf{z}_{t+1}, w(t))$ can be seen as an observation of the gradient $\nabla C(w(t))$. Therefore, in our DALVQ algorithm, each processor $i \in \{1, \dots, M\}$ is able to compute such observations using its own data $\mathbf{z}_1^i, \mathbf{z}_2^i, \dots$. Thus, the DALVQ procedure is defined by equation (3.2) with the following choice for the descent term s^i :

$$s^i(t) = \begin{cases} -\varepsilon_{t+1}^i H(\mathbf{z}_{t+1}^i, w^i(t)) & \text{if } t \in T^i; \\ 0 & \text{otherwise;} \end{cases} \quad (4.2)$$

where $\{\varepsilon_t^i\}_{t=0}^\infty$ are $(0, 1)$ -valued sequences. The sets T^i contain the time instants where the version w^i , kept by processor i , is updated with the descent terms. This fine grain description of the algorithm allows some processors to be idle for computing descent terms (when $t \notin T^i$). This reflects the fact that the computing operations might not take the same time for all processors, which is precisely the core of asynchronous algorithms analysis. Similarly to time delays and combining coefficients, the sets T^i are supposed to be

deterministic but do not need to be known *a priori* for the execution of the algorithm.

In the DALVQ model, randomness arises from the data \mathbf{z} . Therefore, it is natural to let $\{\mathcal{F}_t\}_{t=0}^\infty$ be the filtration built on the σ -algebras

$$\mathcal{F}_t \triangleq \sigma(\mathbf{z}_s^i, i \in \{1, \dots, M\} \text{ and } t \geq s \geq 0), \quad t \geq 0.$$

An easy verification shows that, for all $j \in \{1, \dots, M\}$ and $t \geq 0$, $w^*(t)$ and $w^j(t)$ are \mathcal{F}_t -measurable random variables.

For simplicity, the assumption on the decreasing speed of the sequences $\{\varepsilon_t^i\}_{t=0}^\infty$ is strengthened as follows. The notation $a \vee b$ stands for the maximum of two reals a and b .

Assumption 4.1 *There exist two real numbers $K_1 > 0$ and $K_2 \geq 1$ such that*

$$\frac{K_1}{t \vee 1} \leq \varepsilon_{t+1}^i \leq \frac{K_2}{t \vee 1}, \quad i \in \{1, \dots, M\} \text{ and } t \geq 0.$$

If Assumption 4.1 holds then, the sequences $\{\varepsilon_t^i\}_{t=0}^\infty$ satisfy the standard Assumption 2.2 for stochastic optimization algorithms. Note that the choice of steps proportional to $1/t$ has been proved to be a satisfactory learning rate, theoretically speaking and also for practical implementations (see for instance Murata [28] and Bottou and LeCun [12]).

For practical implementation, the sequences $\{\varepsilon_{t+1}^i\}_{t=0}^\infty$ satisfying Assumption 4.1 can be implemented without a global clock, that is, without assuming that the current value of t is known by the agents. This assumption is satisfied, for example, by taking the current value of ε_t^i proportional to $1/n_t^i$, where n_t^i is the number of times that processor i as performed an update, i.e., the cardinal of the set $T^i \cap \{0, \dots, t\}$. For a given processor, if the time span between consecutive updates is bounded from above and from below, a straightforward examination shows that the sequence of steps satisfy Assumption 4.1.

Finally, the next assumption is essentially technical in nature. It enables to avoid time instants where all processors are idle. It basically requires that, at any time $t \geq 0$, there is at least one processor $i \in \{1, \dots, M\}$ satisfying $s^i(t) \neq 0$.

Assumption 4.2 *One has $\sum_{j=1}^M \mathbb{1}_{\{t \in T^j\}} > 0$.*

4.2 The asynchronous G-Lemma

The aim of this subsection is to state a useful theorem similar to Theorem 2.1, but adapted to our asynchronous distributed context. The precise Definition 3.2 of the agreement vector sequence should not cast aside the intuitive definition. The reader should keep in mind that the vector $w^*(t)$ is also the asymptotical consensus if descent terms are zero after time t . Consequently, even if the agreement vector $\{w^*(t)\}_{t=0}^\infty$ is adapted to the filtration $\{\mathcal{F}_t\}_{t=0}^\infty$, the vector $w^*(t)$ cannot be accessible for a user at time t . Nevertheless, the agreement vector $w^*(t)$ can be interpreted as a ‘‘probabilistic state’’ of the whole distributed quantization scheme at time t . This explains why the agreement vector is a such convenient tool for the analysis of the DALVQ convergence and will be central in our adaptation of G-Lemma, Theorem 4.1.

Let us remark that equation (3.5) writes, for all $t \geq 0$,

$$\begin{aligned} w^*(t+1) &= w^*(t) + \sum_{j=1}^M \phi^j(t) s^j(t) \\ &= w^*(t) - \sum_{j=1}^M \mathbf{1}_{\{t \in T^j\}} \phi^j(t) \varepsilon_{t+1}^j H(\mathbf{z}_{t+1}^j, w^j(t)). \end{aligned}$$

Using the function h defined by identity (2.5) and the fact that the random variables $w^*(t)$ and $w^j(t)$ are \mathcal{F}_t -measurable then, it holds

$$h(w^*(t)) = \mathbb{E} \{ H(\mathbf{z}_{t+1}, w^*(t)) \mid \mathcal{F}_t \}, \quad t \geq 0.$$

and

$$h(w^j(t)) = \mathbb{E} \{ H(\mathbf{z}_{t+1}, w^j(t)) \mid \mathcal{F}_t \}, \quad j \in \{1, \dots, M\} \text{ and } t \geq 0.$$

For all $t \geq 0$, set

$$\varepsilon_{t+1}^* \triangleq \sum_{j=1}^M \mathbf{1}_{\{t \in T^j\}} \phi^j(t) \varepsilon_{t+1}^j. \quad (4.3)$$

Clearly, the real numbers ε_t^* are nonnegative. Their strictly positiveness will be discussed in Proposition 4.1.

Set

$$\Delta M_t^{(1)} \triangleq \sum_{j=1}^M \mathbf{1}_{\{t \in T^j\}} \phi^j(t) \varepsilon_{t+1}^j (h(w^*(t)) - h(w^j(t))), \quad t \geq 0, \quad (4.4)$$

and

$$\Delta M_t^{(2)} \triangleq \sum_{j=1}^M \mathbb{1}_{\{t \in T^j\}} \phi^j(t) \varepsilon_{t+1}^j (h(w^j(t)) - H(\mathbf{z}_{t+1}^j, w^j(t))), \quad t \geq 0. \quad (4.5)$$

Note that $\mathbb{E} \left\{ \Delta M_t^{(2)} \right\} = 0$ and, consequently, that the random variables $\Delta M_t^{(2)}$ can be seen as the increments of a martingale with respect to the filtration $\{\mathcal{F}_t\}_{t=0}^\infty$.

Finally, with this notation, equation (3.5) takes the form

$$w^*(t+1) = w^*(t) - \varepsilon_{t+1}^* h(w^*(t)) + \Delta M_t^{(1)} + \Delta M_t^{(2)}, \quad t \geq 0. \quad (4.6)$$

We are now in a position to state our most useful tool, which is similar in spirit to the G-Lemma, but adapted to the context of distributed asynchronous stochastic gradient descent algorithm.

Theorem 4.1 (Asynchronous G-Lemma) *Assume that the following conditions are satisfied:*

1. $\sum_{t=0}^\infty \varepsilon_t^* = \infty$ and $\varepsilon_t^* \xrightarrow[t \rightarrow \infty]{} 0$.
2. The sequences $\{w^*(t)\}_{t=0}^\infty$ and $\{h(w^*(t))\}_{t=0}^\infty$ are bounded a.s.
3. The series $\sum_{t=0}^\infty \Delta M_t^{(1)}$ and $\sum_{t=0}^\infty \Delta M_t^{(2)}$ converge a.s. in $(\mathbb{R}^d)^\kappa$.
4. There exists a lower semi-continuous function $G : (\mathbb{R}^d)^\kappa \rightarrow [0, \infty)$ such that

$$\sum_{t=0}^\infty \varepsilon_{t+1}^* G(w^*(t)) < \infty, \quad a.s.$$

Then, there exists a connected component Ξ of $\{G = 0\}$ such that

$$\text{dist}(w^*(t), \Xi) \xrightarrow[t \rightarrow \infty]{} 0, \quad a.s.$$

4.3 Trajectory analysis

The Pagès's proof on the almost sure convergence of the CLVQ procedure required a careful examination of the trajectories of the process $\{w(t)\}_{t=0}^\infty$. Thus, in this subsection we investigate similar properties and introduce the assumptions that will be needed to prove our main convergence result, Theorem 4.2.

The next Assumption 4.3 ensures that, for each processor, the quantizers stay in the support of the density.

Assumption 4.3 *One has*

$$\mathbb{P}\{w^j(t) \in \mathcal{G}^\kappa\} = 1, \quad j \in \{1, \dots, M\} \text{ and } t \geq 0.$$

Firstly, let us mention that since the set \mathcal{G}^κ is convex, if Assumption 4.3 holds then,

$$\mathbb{P}\{w^*(t) \in \mathcal{G}^\kappa\} = 1, \quad t \geq 0.$$

Secondly, note that the Assumption 4.3 is not particularly restrictive. This assumption is satisfied under the condition: for each processor, no descent term is added while a combining computation is performed. This writes

$$a_{i,j}(t) = \delta_{i,j} \text{ and } \tau^{i,i}(t) = t, \quad (i, j) \in \{1, \dots, M\}^2 \text{ and } t \in T^i.$$

This requirement makes sense for practical implementations.

Remind that if $t \notin T^i$, then $s^i(t) = 0$. Thus, equation (3.2) takes the form

$$w^i(t+1) = \begin{cases} w^i(t+1) = w^i(t) - \varepsilon_{t+1}^i (w^i(t) - \mathbf{z}_{t+1}^i) & \text{if } t \in T^i; \\ = (1 - \varepsilon_{t+1}^i) w^i(t) + \varepsilon_{t+1}^i \mathbf{z}_{t+1}^i & \\ w^i(t+1) = \sum_{j=1}^M a^{i,j}(t) w^j(\tau^{i,j}(t)) & \text{otherwise.} \end{cases} \quad (4.7)$$

Since \mathcal{G}^κ is a convex set, it follows easily that if $w^j(0) \in \mathcal{G}^\kappa$, then $w^j(t) \in \mathcal{G}^\kappa$ for all $j \in \{1, \dots, M\}$ and $t \geq 0$ and, consequently, that Assumption 4.3 holds.

The next Lemma 4.1 provides a deterministic upper bound on the differences between the distributed versions w^i and the agreement vector. For any subset A of $(\mathbb{R}^d)^\kappa$, the notation $\text{diam}(A)$ stands for the usual diameter defined by

$$\text{diam}(A) = \sup_{x, y \in A} \{\|x - y\|\}.$$

Lemma 4.1 *Assume that Assumptions 4.1 and 4.3 are satisfied then,*

$$\|w^*(t) - w^i(t)\| \leq \sqrt{\kappa} M \text{diam}(\mathcal{G}) A K_2 \theta_t, \quad i \in \{1, \dots, M\} \text{ and } t \geq 0, \text{ a.s.},$$

where $\theta_t \triangleq \sum_{\tau=-1}^{t-1} \frac{1}{\tau \vee 1} \rho^{t-\tau}$ and A (resp. K_2) is the constant introduced in Lemma 3.2 (resp. Assumption 4.1).

The sequence $\{\theta_t\}_{t=0}^\infty$ defined in Lemma 4.1 satisfies

$$\theta_t \xrightarrow[t \rightarrow \infty]{} 0 \text{ and } \sum_{t=0}^{\infty} \frac{\theta_t}{t} < \infty. \quad (4.8)$$

Thus, under Assumptions 4.1 and 4.3, it follows easily that

$$w^*(t) - w^i(t) \xrightarrow[t \rightarrow \infty]{} 0, \quad i \in \{1, \dots, M\}, \text{ a.s.},$$

and

$$w^i(t) - w^j(t) \xrightarrow[t \rightarrow \infty]{} 0, \quad (i, j) \in \{1, \dots, M\}^2, \text{ a.s.} \quad (4.9)$$

This shows that the trajectories of the distributed versions of the quantizers reach asymptotically a consensus with probability 1. In other words, if one of the sequences $\{w^i(t)\}_{t=0}^{\infty}$ converges then, they all converge towards the same value. The rest of the paper is devoted to prove that this common value is in fact a zero of ∇C , i.e. a critical point.

To prove the result mentioned above, we will need the following assumption, which basically states that the components of w^* are parted, for every time t but also asymptotically. This assumption is similar in spirit to the main requirement of Theorem 2.2.

Assumption 4.4 *One has*

1. $\mathbb{P} \{w^*(t) \in \mathcal{D}_*^{\kappa}\} = 1, \quad t \geq 0.$
2. $\mathbb{P} \{\liminf_{t \rightarrow \infty} \text{dist}(w^*(t), \mathcal{C}\mathcal{D}_*^{\kappa}) > 0\} = 1, \quad t \geq 0.$

4.4 Consistency of the DALVQ

In this subsection we state our main theorem on the consistency of the DALVQ. Its proof is based on the Asynchronous G-Lemma, Theorem 4.1. The goal of the next proposition is to ensure that the first assumption of Theorem 4.1 holds.

Proposition 4.1 *Let Assumptions 4.1 and 4.2 hold then, $\varepsilon_t^* > 0, t \geq 0$, $\varepsilon_t^* \xrightarrow[t \rightarrow \infty]{} 0$ and $\sum_{t=0}^{\infty} \varepsilon_t^* = \infty$.*

The second condition required in Theorem 4.1 deals with the convergence of the two series defined by equations (4.4) and (4.5). The next Proposition 4.2 provides sufficient condition for the almost sure convergence of these series.

Proposition 4.2 *If Assumptions 4.1, 4.3 and 4.4 hold then, the series $\sum_{t=0}^{\infty} \Delta M_t^{(1)}$ and $\sum_{t=0}^{\infty} \Delta M_t^{(2)}$ converge almost surely in $(\mathbb{R}^d)^{\kappa}$.*

This next proposition may be considered has the most important step in the proof of the convergence of the DALVQ. It establishes the convergence of a series of the form $\sum_{t=0}^{\infty} \varepsilon_{t+1} \|\nabla C(w(t))\|^2$. The analysis of the convergence of this type of series is standard for the analysis of stochastic gradient method (see for instance Benveniste et al. [5] and Bottou [9]). In our context, we pursue the fruitful use of the agreement vector sequence, $\{w^*(t)\}_{t=0}^{\infty}$, and its related “steps”, $\{\varepsilon_t^*\}_{t=0}^{\infty}$.

Note that under Assumption 4.4, we have $h(w^*(t)) = \nabla C(w^*(t))$ for all $t \geq 0$, almost surely, therefore the sequence $\{\nabla C(w^*(t))\}_{t=0}^{\infty}$ below is well defined.

Proposition 4.3 *If Assumptions 4.1, 4.3 and 4.4 hold then,*

1. $C(w^*(t)) \xrightarrow[t \rightarrow \infty]{} C_{\infty}$, *a.s.*, where $C_{\infty} \in [0, \infty)$,

- 2.

$$\sum_{t=0}^{\infty} \varepsilon_{t+1}^* \|\nabla C(w^*(t))\|^2 < \infty, \quad a.s. \quad (4.10)$$

Remark that from the convergence of the series given by equation (4.10) one can only deduce that $\liminf_{t \rightarrow \infty} \|\nabla C(w^*(t))\| = 0$.

We are now in a position to state the main theorem of this paper, which expresses the convergence of the distributed version towards some zero of the gradient of the distortion. In addition, the convergence results (4.9) imply that if a version converges then, all the versions converge towards this value.

Theorem 4.2 (Asynchronous Theorem) *If Assumptions 4.1, 4.2, 4.3 and 4.4 hold then,*

$$w^*(t) \xrightarrow[t \rightarrow \infty]{} \Xi_{\infty}, \quad a.s.$$

and

$$w^i(t) \xrightarrow[t \rightarrow \infty]{} \Xi_{\infty}, \quad i \in \{1, \dots, M\}, \quad a.s.,$$

where Ξ_{∞} is some connected component of the set $\{\nabla C = 0\} \cap \mathcal{G}^k$.

4.5 Annex

Sketch of the proof of Asynchronous G-Lemma 4.1. The proof is an adaptation of the one found by Fort and Pagès, Theorem 4 in [17]. The

recursive equation (4.6) satisfied by the sequence $\{w^*(t)\}_{t=0}^\infty$ is similar to the iterations (2) in [17] (with the notation of this paper):

$$X^{t+1} = X^t - \varepsilon_{t+1} h(X^t) + \varepsilon_{t+1} (\Delta M^{t+1} + \eta^{t+1}), \quad t \geq 0.$$

Thus, similarly, we define a family of continuous time stepwise function $\{u \mapsto \check{w}(t, u)\}_{t=1}^\infty$.

$$\check{w}^*(0, u) \triangleq w^*(s), \text{ if } u \in [\varepsilon_1^* + \dots + \varepsilon_s^*, \varepsilon_1^* + \dots + \varepsilon_{s+1}^*), \quad u \in [0, \infty).$$

and if $u < \varepsilon_1^*$, $\check{w}^*(0, u) = w^*(0)$.

$$\check{w}^*(t, u) \triangleq \check{w}^*(0, \varepsilon_1^* + \dots + \varepsilon_t^* + u), \quad t \geq 1 \text{ and } u \in [0, \infty).$$

Hence, for every $t \in \mathbb{N}$,

$$\check{w}^*(t, u) = \check{w}^*(0, t) - \int_0^u h(\check{w}^*(t, v)) dv + R_u(t), \quad u \in [0, \infty),$$

where, for every $t \geq 1$ and $u \in [\varepsilon_1^* + \dots + \varepsilon_{t+t'}^*, \varepsilon_1^* + \dots + \varepsilon_{t+t'+1}^*)$,

$$R_u(t) \triangleq \int_{\varepsilon_t^* + \dots + \varepsilon_{t+t'}^*}^{\varepsilon_1^* + \dots + \varepsilon_t^* + u} \check{w}^*(0, v) dv + \sum_{s=t+1}^{t+t'} (\Delta M_s^{(1)} + \Delta M_s^{(2)}).$$

The only difference between the families of continuous time functions $\{\check{w}(t, u)\}_{t=1}^\infty$ and $\{X^{(t)}\}_{t=1}^\infty$ defined in [17] is the remainder term $R_u(t)$. The convergence

$$\sup_{u \in [0, T]} \|R_u(t)\| \xrightarrow{t \rightarrow \infty} 0, \quad T > 0.$$

follows easily from the third assumption of Theorem 4.1. The rest of the proof follows similarly as in Theorem 4 [17].

□

Proof of Lemma 4.1. For all $i \in \{1, \dots, M\}$, and all $t \geq 0$, and all $1 \leq \ell \leq \kappa$, we may write

$$\begin{aligned} & \|w_\ell^i(t) - w_\ell^i(t)\| \\ &= \left\| \sum_{j=1}^M \left((\phi^{i,j}(t, -1) - \phi^j(-1)) w_\ell^j(0) + \sum_{\tau=0}^{t-1} (\phi^{i,j}(t, \tau) - \phi^j(t)) s_\ell^j(\tau) \right) \right\| \\ & \quad \text{(by Definition 3.2 and Lemma 3.1)} \\ &\leq \sum_{j=1}^M |\phi^{i,j}(t, -1) - \phi^j(-1)| \|w_\ell^j(0)\| + \sum_{\tau=0}^{t-1} \sum_{j=1}^M |\phi^{i,j}(t, \tau) - \phi^j(t)| \|s_\ell^j(\tau)\| \\ &\leq A \rho^{t+1} \sum_{j=1}^M \|w_\ell^j(0)\| + A \sum_{\tau=0}^{t-1} \sum_{j=1}^M \rho^{t-\tau} \|s_\ell^j(\tau)\| \\ & \quad \text{(by Lemma 3.2)}. \end{aligned}$$

Thus,

$$\begin{aligned}
& \|w_\ell^i(t) - w_\ell^*(t)\| \\
& \leq A\rho^{t+1} \sum_{j=1}^M \|w_\ell^j(0)\| + A \sum_{\tau=0}^{t-1} \sum_{j=1}^M \rho^{t-\tau} |\varepsilon_{\tau+1}^j| \mathbb{1}_{\{\tau \in T^j\}} \|H(\mathbf{z}_{\tau+1}^j, w^j(\tau))_\ell\| \\
& \quad \text{(by equation (4.2))} \\
& \leq A\rho^{t+1} \sum_{j=1}^M \|w_\ell^j(0)\| \\
& \quad + A \sum_{\tau=0}^{t-1} \sum_{j=1}^M \rho^{t-\tau} |\varepsilon_{\tau+1}^j| \mathbb{1}_{\tau \in T^j} \mathbb{1}_{\{\mathbf{z}_{\tau+1}^j \in W_\ell(w^j(\tau))\}} \|w_\ell^j(\tau) - \mathbf{z}_{\tau+1}^j\|.
\end{aligned}$$

Therefore,

$$\begin{aligned}
& \|w_\ell^i(t) - w_\ell^*(t)\| \\
& \leq AM \operatorname{diam}(\mathcal{G}) \rho^{t+1} + A \operatorname{diam}(\mathcal{G}) K_2 M \sum_{\tau=0}^{t-1} \frac{1}{\tau \vee 1} \rho^{t-\tau} \\
& \quad \text{(because } 0 \in \mathcal{G} \text{ and by Assumptions 4.1 and 4.3)} \\
& \leq A \operatorname{diam}(\mathcal{G}) K_2 M \sum_{\tau=-1}^{t-1} \frac{1}{\tau \vee 1} \rho^{t-\tau}.
\end{aligned}$$

Consequently,

$$\begin{aligned}
& \|w^*(t) - w^i(t)\| \\
& = \sqrt{\sum_{\ell=1}^{\kappa} \|w_\ell^i(t) - w_\ell^*(t)\|^2} \\
& \leq \sqrt{\kappa} M \operatorname{diam}(\mathcal{G}) AK_2 \sum_{\tau=-1}^{t-1} \frac{1}{\tau \vee 1} \rho^{t-\tau}.
\end{aligned}$$

This proves the desired result. □

Let us now introduce the following events: for any $\delta > 0$ and $t \geq 0$,

$$A_\delta^t \triangleq \{w^*(\tau) \in \mathcal{G}_\delta^\kappa, t \geq \tau \geq 0\}.$$

Remind that the $\mathcal{G}_\delta^\kappa$ is a compact subset of \mathcal{G}^κ defined by equality (2.8). The next lemma establishes a detailed analysis of security regions for the parted components of the sequences $\{w^*(t)\}_{t=0}^\infty$ and $\{w^j(t)\}_{t=0}^\infty$.

Lemma 4.2 *Let Assumptions 4.1 and 4.3 hold. Then,*

1. *there exists an integer $t_\delta^1 \geq 0$ such that*

$$A_\delta^t \subset A_{\delta/2}^{t+1}, \quad t \geq t_\delta^1.$$

Moreover,

$$w^*(t) \in \mathcal{G}_\delta^\kappa \Rightarrow [w^*(t), w^*(t+1)] \subset \mathcal{G}_{\delta/2}^\kappa, \quad t \geq t_\delta^1.$$

2. *There exists an integer $t_\delta^2 \geq 0$ such that*

$$w^*(t) \in \mathcal{G}_\delta^\kappa \Rightarrow [w^*(t), w^i(t)] \subset \mathcal{G}_{\delta/2}^\kappa, \quad i \in \{1, \dots, M\} \text{ and } t \geq t_\delta^2.$$

Proof of Lemma 4.2. Proof of statement 1. The proof starts with the observation that under Assumption 4.3 we have $w^j(t) \in \mathcal{G}^\kappa$, for all $i \in \{1, \dots, M\}$ and $t \geq 0$. It follows that, for any $1 \leq \ell \leq \kappa$,

$$\begin{aligned} \|H(\mathbf{z}_{t+1}^j, w^j(t))_\ell\| &\leq \|\mathbf{z}_{t+1}^j - w_\ell^j(t)\| \\ &\leq \text{diam}(\mathcal{G}). \end{aligned}$$

Let us now provide an upper bound on the norm of the differences between two consecutive values of the agreement vector sequence. We may write, for all $t \geq 0$ and all $1 \leq \ell \leq M$,

$$\begin{aligned} &\|w_\ell^*(t+1) - w_\ell^*(t)\| \\ &= \left\| \sum_{j=1}^M \phi^j(t) s_\ell^j(t) \right\| \\ &\leq \sum_{j=1}^M \phi^j(t) \|s_\ell^j(t)\| \\ &\leq \sum_{j=1}^M \varepsilon_{t+1}^j \mathbf{1}_{\{t \in T^j\}} \|H(\mathbf{z}_{t+1}^j, w^j(t))_\ell\| \\ &\quad \text{(by equation (4.2) and statement 1. of Lemma 3.2)} \\ &\leq \frac{M \text{diam}(\mathcal{G}) K_2}{t \vee 1} \\ &\quad \text{(by Assumption 4.1)}. \end{aligned} \tag{4.11}$$

Take $t \geq \frac{4}{\delta} M \text{diam}(\mathcal{G}) K_2$ and $1 \leq k \neq \ell \leq M$. Let α be a real number in the interval $[0, 1]$.

If $w^\star(t) \in \mathcal{G}_\delta^\kappa$ then,

$$\begin{aligned}
& \|(1 - \alpha)w_\ell^\star(t) + \alpha w_\ell^\star(t + 1) - (1 - \alpha)w_k^\star(t) - \alpha w_k^\star(t + 1)\| \\
& \geq \|w_\ell^\star(t) - w_k^\star(t) + \alpha(w_\ell^\star(t + 1) - w_\ell^\star(t)) + \alpha(w_k^\star(t) - w_k^\star(t + 1))\| \\
& \geq \|w_\ell^\star(t) - w_k^\star(t)\| - \|\alpha(w_\ell^\star(t + 1) - w_\ell^\star(t)) + \alpha(w_k^\star(t) - w_k^\star(t + 1))\| \\
& \geq \|w_\ell^\star(t) - w_k^\star(t)\| - \alpha \|w_\ell^\star(t + 1) - w_\ell^\star(t)\| - \alpha \|w_k^\star(t) - w_k^\star(t + 1)\| \\
& \geq \delta - 2\alpha \frac{\delta}{4} \\
& \geq \delta/2.
\end{aligned}$$

This proves that the whole segment $[w^\star(t), w^\star(t + 1)]$ is contained in $\mathcal{G}_{\delta/2}^\kappa$.

Proof of statement 2. Take $t \geq 0$ and $1 \leq \ell \leq M$. If $w^\star(t) \in \mathcal{G}_\delta^\kappa$ then, by Lemma 4.1, there exists t_δ^2 such that

$$\|w_\ell^\star(t) - w_\ell^i(t)\| \leq \frac{\delta}{4}, \quad i \in \{1, \dots, M\} \text{ and } t \geq t_\delta^2.$$

Let k and ℓ two distinct integers between 1 and M . For any $t \geq t_\delta^2$,

$$\begin{aligned}
& \|\alpha w_k^i(t) + (1 - \alpha)w_k^\star(t) - \alpha w_\ell^i(t) - (1 - \alpha)w_\ell^\star(t)\| \\
& \geq \|w_k^\star(t) - w_\ell^\star(t) + \alpha(w_k^i(t) - w_k^\star(t)) + \alpha(w_\ell^\star(t) - w_\ell^i(t))\| \\
& \geq \|w_k^\star(t) - w_\ell^\star(t)\| - \alpha \|w_k^i(t) - w_k^\star(t)\| - \alpha \|w_\ell^\star(t) - w_\ell^i(t)\| \\
& \geq \delta - 2\alpha \frac{\delta}{4} \\
& \geq \delta/2.
\end{aligned}$$

This implies $[w^\star(t), w^i(t)] \subset \mathcal{G}_{\delta/2}^\kappa$, as desired. □

Proof of Proposition 4.1. By definition ε_{t+1}^\star equals $\sum_{j=1}^M \mathbf{1}_{\{t \in T^j\}} \phi^j(t) \varepsilon_{t+1}^j$, for all $t \geq 0$.

On the one hand, since the real number $\phi^j(t)$ belongs to the interval $[\eta, 1]$ (by Lemma 3.2) ε_{t+1}^\star is bounded from above by $\frac{MK_2}{tV_1}$ using the right-hand side inequality of Assumption 4.1.

On the other hand, ε_{t+1}^\star is bounded from below by the nonnegative real number $\eta \frac{K_1}{tV_1} \sum_{j=1}^M \mathbf{1}_{\{t \in T^j\}}$ using the left-hand side inequality of Assumption

4.1. Note also that as Assumption 4.2 holds, this real number is a positive one. Therefore, it follows that

$$\varepsilon_t^* \xrightarrow[t \rightarrow \infty]{} 0$$

and

$$\sum_{t=0}^{\infty} \varepsilon_t^* = \infty.$$

□

Proof of Proposition 4.2. Consistency of $\sum_{t=0}^{\infty} \Delta M_t^{(1)}$. Let δ be a positive real number and let $t \geq t_\delta^2$, where t_δ^2 is given by Lemma 4.10. We may write

$$\begin{aligned} & \mathbb{1}_{A_\delta^t} \sum_{j=1}^M \mathbb{1}_{\{t \in T^j\}} \phi^j(t) \varepsilon_{t+1}^j \|h(w^*(t)) - h(w^j(t))\| \\ & \leq \mathbb{1}_{\{[w^*(t), w^j(t)] \subset \mathcal{G}_{\delta/2}^\kappa\}} \sum_{j=1}^M \phi^j(t) \varepsilon_{t+1}^j \|\nabla C(w^*(t)) - \nabla C(w^j(t))\| \\ & \quad \text{(using statement 2. of Lemma 4.2 and the fact that } \nabla C = h \text{ on } \mathcal{D}_*^\kappa) \\ & \leq \mathbb{1}_{\{[w^*(t), w^j(t)] \subset \mathcal{G}_{\delta/2}^\kappa\}} P_{\delta/2} \sum_{j=1}^M \varepsilon_{t+1}^j \|w^*(t) - w^j(t)\| \\ & \quad \text{(by Lemma 2.1)} \\ & \leq \sqrt{\kappa} \text{diam}(\mathcal{G}) AK_2^2 P_{\delta/2} M \frac{\theta_t}{t} \\ & \quad \text{(by Lemma 4.1).} \end{aligned}$$

Thus, since $\sum_{t=0}^{\infty} \frac{\theta_t}{t} < \infty$, the series

$$\sum_{t=0}^{\infty} \mathbb{1}_{A_\delta^t} \sum_{j=1}^M \mathbb{1}_{\{t \in T^j\}} \phi^j(t) \varepsilon_{t+1}^j \|h(w^*(t)) - h(w^j(t))\|$$

is almost surely convergent. Under Assumption 4.4, we have

$$\mathbb{P} \left\{ \bigcup_{\delta > 0} \bigcap_{t \geq 0} A_\delta^t \right\} = 1.$$

It follows that the series $\sum_{t=0}^{\infty} \Delta M_t^{(1)}$ converges almost surely in $(\mathbb{R}^d)^\kappa$.

Consistency of $\sum_{t=0}^{\infty} \Delta M_t^{(2)}$. The sequence of random variables $M_t^{(2)}$ defined, for all $t \geq 0$, by

$$\begin{aligned} M_t^{(2)} &\triangleq \sum_{\tau=0}^t \Delta M_{\tau}^{(2)} \\ &= \sum_{\tau=0}^t \sum_{j=1}^M \mathbf{1}_{\{\tau \in T^j\}} \varepsilon_{\tau+1}^j \phi^j(\tau) (h(w^j(\tau)) - H(\mathbf{z}_{\tau+1}^j, w^j(\tau))). \end{aligned}$$

is a vector valued martingale with respect to the filtration $\{\mathcal{F}_t\}_{t=0}^{\infty}$. It turns out that this martingale has square integrable increments. Precisely,

$$\sum_{t=0}^{\infty} \mathbb{E} \left\{ \left\| M_{t+1}^{(2)} - M_t^{(2)} \right\|^2 \mid \mathcal{F}_t \right\} = \sum_{t=1}^{\infty} \mathbb{E} \left\{ \left\| \Delta M_t^{(2)} \right\|^2 \mid \mathcal{F}_t \right\} < \infty.$$

Indeed, for all $j \in \{1, \dots, M\}$ and $t \geq 1$,

$$\begin{aligned} &\sum_{\tau=1}^t \mathbb{E} \left\{ \left\| \mathbf{1}_{\{\tau \in T^j\}} \varepsilon_{\tau+1}^j (h(w^j(\tau)) - H(\mathbf{z}_{\tau+1}^j(\tau), w^j(\tau))) \right\|^2 \mid \mathcal{F}_{\tau} \right\} \\ &\leq \sum_{\tau=1}^t (\varepsilon_{\tau+1}^j)^2 \mathbb{E} \left\{ \left\| h(w^j(\tau)) - H(\mathbf{z}_{\tau+1}^j(\tau), w^j(\tau)) \right\|^2 \mid \mathcal{F}_{\tau} \right\} \\ &\leq 2 \sum_{\tau=1}^t (\varepsilon_{\tau+1}^j)^2 \mathbb{E} \left\{ \left\| h(w^j(\tau)) \right\|^2 + \left\| H(\mathbf{z}_{\tau+1}^j(\tau), w^j(\tau)) \right\|^2 \mid \mathcal{F}_{\tau} \right\} \\ &\leq 2\kappa \text{diam}(\mathcal{G})^2 \sum_{\tau=1}^t (\varepsilon_{\tau+1}^j)^2 \\ &\quad \text{(using Assumption 4.3)} \\ &\leq 2\kappa \text{diam}(\mathcal{G})^2 K_2^2 \sum_{\tau=1}^t \frac{1}{\tau^2}. \end{aligned}$$

We conclude that the series $\sum_{t \geq 1} \Delta M_t^{(2)}$ is almost surely convergent. □

Proof of proposition 4.3. Denote by $\langle x, y \rangle$ the canonical inner product of two vectors $x, y \in \mathbb{R}^d$ and also, with a slight abuse of notation, the canonical inner product of two vectors $x, y \in (\mathbb{R}^d)^{\kappa}$. Let δ be a positive real number. Take any $t \geq \max\{t_{\delta}^1, t_{\delta}^2\}$, where t_{δ}^1 and t_{δ}^2 are defined as in Lemma 4.2. For

simplicity we assume that $\max\{t_\delta^1, t_\delta^2\} \geq 1$. One has,

$$\begin{aligned}
& \mathbb{1}_{A_\delta^{t+1}} C(w^*(t+1)) \\
& \leq \mathbb{1}_{A_\delta^t} C(w^*(t+1)) \\
& \quad \text{(by definition } A_\delta^{t+1} \subset A_\delta^t \text{)} \\
& \leq \mathbb{1}_{A_\delta^t} C(w^*(t)) + \mathbb{1}_{A_\delta^t} \langle \nabla C(w^*(t)), w^*(t+1) - w^*(t) \rangle \\
& \quad + \mathbb{1}_{\{[w^*(t), w^*(t+1)] \subset \mathcal{G}_{\delta/2}^k\}} \\
& \quad \times \left[\sup_{z \in [w^*(t), w^*(t+1)]} \{ \|\nabla C(z) - \nabla C(w^*(t))\| \} \|w^*(t+1) - w^*(t)\| \right] \\
& \leq \mathbb{1}_{A_\delta^t} C(w^*(t)) + \mathbb{1}_{A_\delta^t} \langle \nabla C(w^*(t)), w^*(t+1) - w^*(t) \rangle \\
& \quad + P_{\delta/2} \|w^*(t+1) - w^*(t)\|^2 \\
& \quad \text{(using Lemma 2.1.)}
\end{aligned}$$

The second inequality above holds since the bounded increment formula above is valid by statement 1 of Lemma 4.2.

$$\begin{aligned}
& \mathbb{1}_{A_\delta^{t+1}} C(w^*(t+1)) \\
& \leq \mathbb{1}_{A_\delta^t} C(w^*(t)) + \mathbb{1}_{A_\delta^t} \langle \nabla C(w^*(t)), \sum_{j=1}^M \phi^j(t) s^j(t) \rangle \\
& \quad \text{(by equation (3.5))} \\
& \quad + \kappa P_{\delta/2} \left(\frac{K_2 M \text{diam}(\mathcal{G})}{t} \right)^2 \\
& \leq \mathbb{1}_{A_\delta^t} C(w^*(t)) + \mathbb{1}_{A_\delta^t} \sum_{j=1}^M \langle \nabla C(w^j(t)), \phi^j(t) s^j(t) \rangle \\
& \quad + \mathbb{1}_{A_\delta^t} \sum_{j=1}^M \langle \nabla C(w^*(t)) - \nabla C(w^j(t)), \phi^j(t) s^j(t) \rangle \\
& \quad + \kappa P_{\delta/2} \left(\frac{K_2 M \text{diam}(\mathcal{G})}{t} \right)^2.
\end{aligned}$$

The first inequality above holds by inequality (4.11) provided in the proof of

Lemma 4.2).

$$\begin{aligned}
& \mathbb{1}_{A_\delta^{t+1}} C(w^*(t+1)) \\
& \leq \mathbb{1}_{A_\delta^t} C(w^*(t)) + \mathbb{1}_{A_\delta^t} \sum_{j=1}^M \langle \nabla C(w^j(t)), \phi^j(t) s^j(t) \rangle \\
& \quad + \mathbb{1}_{A_\delta^t} \sum_{j=1}^M |\langle \nabla C(w^*(t)) - \nabla C(w^j(t)), \phi^j(t) s^j(t) \rangle| \\
& \quad + \kappa P_{\delta/2} \left(\frac{K_2 M \text{diam}(\mathcal{G})}{t} \right)^2 \\
& \leq \mathbb{1}_{A_\delta^t} C(w^*(t)) + \mathbb{1}_{A_\delta^t} \sum_{j=1}^M \langle \nabla C(w^j(t)), \phi^j(t) s^j(t) \rangle \\
& \quad + \sum_{j=1}^M \mathbb{1}_{A_\delta^t} \|\nabla C(w^*(t)) - \nabla C(w^j(t))\| \|\phi^j(t) s^j(t)\| \\
& \quad \text{(using Cauchy-Schwarz inequality)} \\
& \quad + \kappa P_{\delta/2} \left(\frac{K_2 M \text{diam}(\mathcal{G})}{t} \right)^2.
\end{aligned}$$

Thus,

$$\begin{aligned}
& \mathbb{1}_{A_\delta^{t+1}} C(w^*(t+1)) \\
& \leq \mathbb{1}_{A_\delta^t} C(w^*(t)) + \mathbb{1}_{A_\delta^t} \sum_{j=1}^M \langle \nabla C(w^j(t)), \phi^j(t) s^j(t) \rangle \\
& \quad + \sum_{j=1}^M \mathbb{1}_{\{[w^*(t), w^j(t)] \subset \mathcal{G}_{\delta/2}^\kappa\}} \|\nabla C(w^*(t)) - \nabla C(w^j(t))\| \|\phi^j(t) s^j(t)\| \\
& \quad \text{(by statement 2 of Lemma 4.2)} \\
& \quad + \kappa P_{\delta/2} \left(\frac{K_2 M \text{diam}(\mathcal{G})}{t} \right)^2
\end{aligned}$$

Therefore,

$$\begin{aligned}
& \mathbb{1}_{A_\delta^{t+1}} C(w^*(t+1)) \\
& \leq \mathbb{1}_{A_\delta^t} C(w^*(t)) + \mathbb{1}_{A_\delta^t} \sum_{j=1}^M \langle \nabla C(w^j(t)), \phi^j(t) s^j(t) \rangle \\
& \quad + P_{\delta/2} \sum_{j=1}^M \|w^*(t) - w^j(t)\| \|\phi^j(t) s^j(t)\| \\
& \quad \text{(using Lemma 2.1)} \\
& \quad + \kappa P_{\delta/2} \left(\frac{K_2 M \text{diam}(\mathcal{G})}{t} \right)^2.
\end{aligned}$$

Hence,

$$\begin{aligned}
& \mathbb{1}_{A_\delta^{t+1}} C(w^*(t+1)) \\
& \leq \mathbb{1}_{A_\delta^t} C(w^*(t)) + \mathbb{1}_{A_\delta^t} \sum_{j=1}^M \langle \nabla C(w^j(t)), \phi^j(t) s^j(t) \rangle \\
& \quad + P_{\delta/2} A K_2^2 \kappa M^2 \text{diam}(\mathcal{G})^2 \frac{\theta_t}{t} \\
& \quad + \kappa P_{\delta/2} \left(\frac{K_2 M \text{diam}(\mathcal{G})}{t} \right)^2. \tag{4.12}
\end{aligned}$$

Set

$$\Omega_\delta^1 \triangleq P_{\delta/2} A K_2^2 \kappa M^2 \text{diam}(\mathcal{G})^2$$

and

$$\Omega_\delta^2 \triangleq P_{\delta/2} (K_2 M \text{diam}(\mathcal{G}))^2.$$

In the sequel, we shall need the following lemma.

Lemma 4.3 *For all $t \geq t \geq \max\{t_\delta^1, t_\delta^2\}$, the quantity W_t (??) below is a nonnegative supermartingale with respect to the filtration $\{\mathcal{F}_t\}_{t=0}^\infty$:*

$$\begin{aligned}
W_t \triangleq & \mathbb{1}_{A_\delta^t} C(w^*(t)) + \frac{\eta K_1}{2} \sum_{\tau=0}^{t-1} \mathbb{1}_{A_\delta^\tau} \frac{1}{\tau} \sum_{j=1}^M \mathbb{1}_{\{\tau \in T^j\}} \|\nabla C(w^j(\tau))\|^2 \\
& + \Omega_\delta^1 \sum_{\tau=t}^\infty \frac{\theta(\tau)}{\tau} + \Omega_\delta^2 \sum_{\tau=t}^\infty \frac{1}{\tau^2}, \quad t \geq 1.
\end{aligned}$$

Proof of Lemma 4.3. Indeed, using the upper bound provided by equation (4.12),

$$\begin{aligned}
& \mathbb{E} \left\{ \mathbf{1}_{A_\delta^{t+1}} C(w^*(t+1)) \mid \mathcal{F}_t \right\} \\
& \leq \mathbf{1}_{A_\delta^t} C(w^*(t)) + \mathbf{1}_{A_\delta^t} \sum_{j=1}^M \mathbb{E} \left\{ \langle \nabla C(w^j(t)), \phi^j(t) s^j(t) \rangle \mid \mathcal{F}_t \right\} \\
& \quad + \Omega_\delta^1 \frac{1}{t} \theta_t + \Omega_\delta^2 \frac{1}{t^2} \\
& \leq \mathbf{1}_{A_\delta^t} C(w^*(t)) \\
& \quad + \mathbf{1}_{A_\delta^t} \sum_{j=1}^M \left\langle \nabla C(w^j(t)), \mathbb{E} \left\{ -\frac{1}{2} \mathbf{1}_{\{t \in T^j\}} \phi^j(t) \varepsilon_{t+1}^j H(\mathbf{z}_{t+1}^j, w^j(t)) \mid \mathcal{F}_t \right\} \right\rangle \\
& \quad + \Omega_\delta^1 \frac{\theta_t}{t} + \Omega_\delta^2 \frac{1}{t^2} \\
& \leq \mathbf{1}_{A_\delta^t} C(w^*(t)) \\
& \quad - \frac{1}{2} \mathbf{1}_{A_\delta^t} \sum_{j=1}^M \mathbf{1}_{\{t \in T^j\}} \phi^j(t) \varepsilon_{t+1}^j \|\nabla C(w^j(t))\|^2 + \Omega_\delta^1 \frac{\theta_t}{t} + \Omega_\delta^2 \frac{1}{t^2} \\
& \leq \mathbf{1}_{A_\delta^t} C(w^*(t)) \\
& \quad - \frac{\eta K_1}{2t} \mathbf{1}_{A_\delta^t} \sum_{j=1}^M \mathbf{1}_{\{t \in T^j\}} \|\nabla C(w^j(t))\|^2 + \Omega_\delta^1 \frac{\theta_t}{t} + \Omega_\delta^2 \frac{1}{t^2}.
\end{aligned}$$

In the last inequality we used the fact that $\phi^j(t) \geq \eta$ (Lemma 3.2) and $\varepsilon_{t+1}^j \geq \frac{K_1}{t}$ (Assumption 4.1).

□

Proof of proposition 4.3 (continued). Since $\{W_t\}_{t=1}^\infty$ is a nonnegative supermartingale, W_t converges almost surely as $t \rightarrow \infty$. Then, as $\sum_{\tau=t}^\infty \frac{\theta(\tau)}{\tau} \xrightarrow[t \rightarrow \infty]{} 0$ and $\sum_{\tau=t}^\infty \frac{1}{\tau^2} \xrightarrow[t \rightarrow \infty]{} 0$, we have

$$\mathbf{1}_{A_\delta^t} C(w^*(t)) \xrightarrow[t \rightarrow \infty]{} C_\infty, \quad \text{a.s.}, \quad (4.13)$$

where $C_\infty \in [0, \infty)$ and

$$\sum_{\tau=0}^\infty \mathbf{1}_{A_\delta^\tau} \frac{1}{\tau \vee 1} \sum_{j=1}^M \mathbf{1}_{\{\tau \in T^j\}} \|\nabla C(w^j(\tau))\|^2 < \infty, \quad \text{a.s.} \quad (4.14)$$

Proof of statement 1. Assumption 4.4 means that

$$\mathbb{P} \left\{ \bigcup_{\delta > 0} \bigcap_{t \geq 0} A_\delta^t \right\} = 1.$$

Statement 1 follows easily from the convergence (4.13).

Proof of statement 2. The required convergence (4.10) is proven as follows. We have

$$\begin{aligned} & \sum_{\tau=0}^t \varepsilon_{\tau+1}^* \mathbf{1}_{A_\delta^\tau} \|\nabla C(w^*(\tau))\|^2 \\ & \leq \sum_{\tau=0}^t \sum_{j=1}^M \phi^j(\tau) \mathbf{1}_{\{\tau \in T^j\}} \mathbf{1}_{A_\delta^\tau} \varepsilon_{\tau+1}^j \|\nabla C(w^*(\tau))\|^2 \\ & \leq 2K_2 \sum_{\tau=0}^t \mathbf{1}_{A_\delta^\tau} \frac{1}{\tau \vee 1} \sum_{j=1}^M \mathbf{1}_{\{\tau \in T^j\}} \|\nabla C(w^j(\tau))\|^2 \\ & \quad (\text{using Assumption 4.2}) \\ & + 2K_2 \sum_{\tau=0}^t \mathbf{1}_{\{[w^*(\tau), w^j(\tau)] \subset \mathcal{G}_{\delta/2}^c\}} \frac{1}{\tau \vee 1} \sum_{j=1}^M \|\nabla C(w^j(\tau)) - \nabla C(w^*(\tau))\|^2 \\ & \quad (\text{using Assumption 4.2 and statement 2 of Lemma 4.2.}) \end{aligned}$$

Thus,

$$\begin{aligned} & \sum_{\tau=0}^t \varepsilon_{\tau+1}^* \mathbf{1}_{A_\delta^\tau} \|\nabla C(w^*(\tau))\|^2 \\ & \leq 2K_2 \sum_{\tau=0}^t \mathbf{1}_{A_\delta^\tau} \frac{1}{\tau \vee 1} \sum_{j=1}^M \mathbf{1}_{\{\tau \in T^j\}} \|\nabla C(w^j(\tau))\|^2 \\ & + 2K_2 P_{\delta/2}^2 \sum_{\tau=0}^t \mathbf{1}_{\{[w^*(\tau), w^j(\tau)] \subset \mathcal{G}_{\delta/2}^c\}} \frac{1}{\tau \vee 1} \sum_{j=1}^M \|w^j(\tau) - w^*(\tau)\|^2 \\ & \quad (\text{by Lemma 2.1}). \end{aligned}$$

Thus,

$$\begin{aligned}
& \sum_{\tau=0}^t \varepsilon_{\tau+1}^* \mathbb{1}_{A_\delta^\tau} \|\nabla C(w^*(\tau))\|^2 \\
& \leq 2K_2 \sum_{\tau=0}^t \mathbb{1}_{A_\delta^\tau} \frac{1}{\tau \vee 1} \sum_{j=1}^M \mathbb{1}_{\{\tau \in T^j\}} \|\nabla C(w^j(\tau))\|^2 \\
& \quad + 2P_{\delta/2}^2 K_2^3 \kappa M^3 A^2 \text{diam}(\mathcal{G})^2 \sum_{\tau=1}^t \frac{1}{\tau \vee 1} \theta_\tau^2 \\
& \quad \text{(by Lemma 4.1)}.
\end{aligned}$$

Finally, using the convergence (4.14), one has

$$\sum_{\tau=0}^{\infty} \varepsilon_{\tau+1}^* \mathbb{1}_{A_\delta^\tau} \|\nabla C(w^*(\tau))\|^2 < \infty, \quad \text{a.s.},$$

and the conclusion follows from the fact that Assumption 4.4 implies

$$\mathbb{P} \left\{ \bigcup_{\delta>0} \bigcap_{t \geq 0} A_\delta^t \right\} = 1.$$

□

Proof of Theorem 4.2. The proof consists in verifying the assumptions of Theorem 4.1 with the function $\widehat{G}(w)$ defined by equation (2.9).

It has been outlined that Assumption 4.3 implies that $w^*(t)$ lie in the compact set \mathcal{G}^κ , almost surely, for all $t \geq 0$. For all $\mathbf{z} \in \mathcal{G}$ and all $t \geq 0$, we have $\|H(\mathbf{z}, w^*(t))\| \leq \sqrt{\kappa} \text{diam}(\mathcal{G})$, almost surely, whereas $\{h(w^*(t))\}_{t=0}^\infty$ satisfies

$$h(w^*(t)) = \mathbb{E} \{ H(\mathbf{z}, w^*(t)) \mid \mathcal{F}_t \}, \quad t \geq 0, \text{ a.s.}$$

Thus, the sequences $\{w^*(t)\}_{t=0}^\infty$ and $\{h(w^*(t))\}_{t=0}^\infty$ are bounded almost surely.

Proposition 4.1 (respectively Proposition 4.2, respectively Proposition 4.3 shows that the first assumption (respectively the second assumption, respectively the fourth assumption) of Theorem 4.1 holds. This concludes the proof of the theorem.

□

References

- [1] A. Antos. Improved minimax bounds on the test and training distortion of empirically designed vector quantizers. *IEEE Transactions on Information Theory*, 51, 2005.
- [2] E.A. Abaya, and G.L. Wise. Convergence of vector quantizers with applications to optimal quantization. *SIAM Journal on Applied Mathematics*, 44, 1984.
- [3] A. Antos, L. Györfi, and A. György. Improved convergence rates in empirical vector quantizer design. *IEEE Transactions on Information Theory*, 51, 2005.
- [4] P.L. Bartlett, G.G. Lugosi, and T. Linder. The minimax distortion redundancy in empirical quantizer design. *IEEE Transactions on Information Theory*, 44, 1998.
- [5] A. Benveniste, M. Métivier, and P. Priouret. *Adaptive algorithms and stochastic approximations*, Springer-Verlag, 1990.
- [6] D.P. Bertsekas, and J.N. Tsitsiklis. *Parallel and distributed computation: numerical methods*, Prentice-Hall, 1989.
- [7] G. Biau, L. Devroye, and G. Lugosi. On the performance of clustering in Hilbert spaces. *IEEE Transactions on Information Theory*, 54, 2008.
- [8] V.D. Blondel, J.M. Hendrickx, A. Olshevsky, and J.N. Tsitsiklis. Convergence in multiagent coordination, consensus, and flocking. *Proceedings of the 44th IEEE Conference on Decision and Control*, 2005.
- [9] L. Bottou. Stochastic gradient learning in neural networks. *Proceedings of Neuro-Nîmes*, 1991.
- [10] L. Bottou. Online algorithms and stochastic approximations. *Online Learning and Neural Networks*. Cambridge University Press, 1998.
- [11] L. Bottou. Convergence properties of the k -means algorithm. *Advances in Neural Information Processing Systems*, 7, 1995.
- [12] L. Bottou, and Y. LeCun. Large scale online learning. *Advances in Neural Information Processing Systems*, 16, 2004.

- [13] F. Bullo, J. Cortés, and S. Martínez. *Distributed control of robotic networks*. 2008. Manuscript preprint. Electronically available at <http://coordinationbook.info>.
- [14] P.A. Chou. The distortion of vector quantizers trained on n vectors decreases to the optimum at $o_p(1/n)$. *IEEE Transactions on Information Theory*, 8, 1994.
- [15] D. Jungnickel. *Graphs, networks and algorithms*, Springer-Verlag, 1999.
- [16] J.C. Fort, and G. Pagès. On the a.s. convergence of the Kohonen algorithm with a general neighborhood function. *The Annals of Applied Probability*, 5, 1995.
- [17] J.C. Fort, and G. Pagès. Convergence of stochastic algorithms: from the Kushner-Clark theorem to the Lyapounov functional method. *Advances in Applied Probability*, 28, 1996.
- [18] A. Gersho, and R.M. Gray. *Vector quantization and signal compression*, Kluwer Academic Publishers, 1992.
- [19] S. Graf, and H. Luschgy. *Foundations of quantization for probability distributions*, Springer-Verlag, 2000.
- [20] M. Inaba, N. Katoh, and H. Imai. Applications of weighted Voronoi diagrams and randomization to variance-based k -clustering: (extended abstract). *Proceedings of the 10th Annual Symposium on Computational Geometry*, 1994.
- [21] T. Kohonen. Analysis of a simple self-organizing process. *Biological Cybernetics*, 44, 1982.
- [22] H.J. Kushner, and D.S. Clark. *Stochastic approximation for constrained and unconstrained systems*, Springer-Verlag, 1978.
- [23] T. Linder. Learning-theoretic methods in vector quantization. *Lecture Notes for the Advanced School on the Principles of Nonparametric Learning*, 2001.
- [24] T. Linder, G. Lugosi, and K. Zeger. Rates of convergence in the source coding theorem, in empirical quantizer design, and in universal lossy source coding. *IEEE Transactions on Information Theory*, 40, 1994.
- [25] S. Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28, 2003.

- [26] J.B. MacQueen. Some methods of classification and analysis of multivariate observations. *Proceedings of the 5h Berkeley Symposium on Mathematical Statistics and Probability*, 1967.
- [27] R. McDonald, K. Hall, and G. Mann. Distributed training strategies for the structured perceptron. *Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 2010.
- [28] N. Murata. A statistical study of on-line learning. *Online Learning and Neural Networks*, Cambridge University Press, 1998.
- [29] G. Pagès. A space vector quantization for numerical integration. *Journal of Applied and Computational Mathematics*, 89, 1997.
- [30] D. Pollard. Quantization and the method of k -means. *IEEE Transactions on Information Theory*, 28, 1982.
- [31] D. Pollard. Strong consistency of k -means clustering. *The Annals of Statistics*, 9, 1981.
- [32] D. Pollard. A central limit theorem for k -means clustering. *The Annals of Probability*, 28, 1982.
- [33] H. Robbins, and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22, 1951.
- [34] M.J. Sabin, and R.M. Gray. Global convergence and empirical consistency of the generalized Lloyd algorithm. *IEEE Transactions on Information Theory*, 32, 1986.
- [35] T. Linder. On the training distortion of vector quantizers. *IEEE Transactions on Information Theory*, 46, 2000.
- [36] J. Tsitsiklis, D. Bertsekas, and M. Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31, 1986.
- [37] J. Tsitsiklis. *Problems in decentralized decision making and computation*. PhD thesis, Department of EECS, MIT, 1984.
- [38] M. Zinkevich, A. Smola, and J. Langford. Slow learners are fast. *Advances in Neural Information Processing Systems*, 22, 2009.