Network Optimization with Heuristic Rational Agents

Ceyhun Eksin and Alejandro Ribeiro
Department of Electrical and Systems Engineering, University of Pennsylvania
{ceksin, aribeiro}@seas.upenn.edu

***Abstract*—We study a distributed model for optimizing a sum of convex objective functions corresponding to agents in the network. At random times, agents execute actions based on heuristic rational rules considering only local information. Heuristic rational rules are probabilistic and their expectation yields the actual optimal action. Under heuristic rational rule iterations, it is shown that global network cost comes within a close vicinity of the optimal value infinitely often with probability 1. Furthermore, an exponential bound on probability of deviating from vicinity of the optimal value is derived. We exemplify heuristic rational behavior on estimation of a random field using a wireless sensor network.**

I. Introduction

Distributed network optimization to model the emergence of global behavior through local actions is receiving attention beyond its use in sensor and communication networks. In, e.g., biological systems, network optimization models that mimic natural phenomena like bird flocking have been introduced [1]–[3]. In social networks, belief propagation and consensus formation [4], [5] can be understood in terms of distributed network optimization. In a fundamental sense, all of these works start from a global objective that the network agents want to optimize – e.g., total drag in a flock of birds – through the selection of local variables – e.g., birds' positions and velocities – while restricting interactions to neighboring agents – e.g., positions and velocities are updated relative to the three closest neighboring birds on the field of vision.

Agents are also assumed to act rationally in that they update their local variables in a manner that is optimal with respect to the available information – e.g., birds choose a position and speed to minimize their own drag. Assuming optimal behavior in the context of natural, as opposed to engineered, networks limits the applicability of these models because making optimal decisions requires exceedingly high levels of awareness and cunningness. The goal of this paper is to study more realistic models whereby agents execute heuristic rational actions that are optimal in an average sense only. Efforts to lift unrealistic assumptions in network optimization include the study of asynchronous updates, time-varying or unreliable communication links [6], [7], or communication contaminated with random noise [1], [8]. Our work differs from these contributions in that we are considering the variable update rules themselves as being imperfect.

We formulate global behavior with a cost given by a sum of local terms involving nonlinear functions of self and neighboring variables. At random times, agents observe the current values of their neighbors' variables and apply a heuristic rule with the intent of minimizing the global cost with respect to the selection of their local variables. These heuristic rules need not be optimal but we assume that they are so in expectation (Section II). Because of the randomness associated with heuristic rational rules we do not expect convergence to optimal global behavior. Consequently, our goal is to describe the difference in the yield of optimal variables and the values achieved by heuristic rational rules (Section III). To characterize this difference, we identify the variable updates with a stochastic coordinate descent algorithm. Exploiting this fact, it is possible to show that for points sufficiently far away from optimal the distance to optimality between subsequent uses of heuristic rational rules satisfies a supermartingale inequality (Lemma 1). This observation leads to the conclusion that a neighborhood of optimality is visited infinitely often with probability 1 (Theorem 1). The size of the near optimality region depends on parameters of the function being minimized and is proportional to the variance of the heuristic rational rule. We further show that between visits to optimality the gap in the yield of agents' variables is bounded exponentially (Section IV). This bound shows that even though it is possible for local variables to become arbitrarily bad, significant deviations are exponentially rare.

We exemplify network optimization using heuristic rational agents with two case studies (Section V). The first example models estimation of temperature on a random field using a network of wireless sensors. Second example considers a network of animals foraging on a field. We conclude the paper in Section VI.

II. Heuristic Rational Decisions

We represent the network of agents by a symmetric graph $G = (V, E)$. Vertices $i \in V$ denote agents and edges $(i, j) \in E$ connections between them. The set of agents that form an edge with agent i is denoted by the neighborhood of agent i, $n(i) = \{j : (j, i) \in E\}$. $N_i := \#(n(i))$ denotes the cardinality of the number of neighbors of agent i. Each of the agents $i \in V$ is associated with corresponding variable $x_i \in \mathbb{R}^n$ and a convex function $f_{0i}(x_i)$. For each edge $(i, j) \in E$, there exists a convex function $f_{ij}(x_i, x_j)$ that depends on the agent variables at the vertices of edge. We require that functions $f_{ij}(x_i, x_j)$ and $f_{ji}(x_j, x_i)$ be equal for all $i, j \in n(i)$ to maintain symmetry. There exists a convex set $\mathcal{X}_i$ that each variable x_i is constrained to; i.e., $x_i \in \mathcal{X}_i \subseteq \mathbb{R}^n$. Further, define $x_{n(i)} = \{x_j\}_{j \in n(i)}$ to denote the variables of all neighbors of i, $x = \{x_i\}_{i \in V}$ to represent all variables, and $\mathcal{X}$ as the Cartesian product of sets $\mathcal{X}_i$.

Work in this paper is supported by ARO P-57920-NS, NSF CAREER CCF-0952867, and NSF CCF-1017454.

We define the incurred cost function for agent i as follows

$$f_i(x_i, x_{n(i)}) := f_{0i}(x_i) + \sum_{j \in n(i)} f_{ij}(x_i, x_j). \quad (1)$$

Agent i would like to minimize (1) by selecting its variable $x_i \in \mathcal{X}_i$ appropriately.

The sum of the local costs $f_i(x_i, x_{n(i)})$ among all agents yields a measure for network aggregate cost $f(x) := \sum_{i \in V} f_i(x_i, x_{n(i)})$. Aggregate cost measures optimality of configuration $x := \{x_i\}_{i \in V}$ from a global perspective. We define an optimal configuration x^* as

$$x^* := \operatorname*{argmin}_{x \in \mathcal{X}} \ f(x) = \operatorname*{argmin}_{x \in \mathcal{X}} \ \sum_{i \in V} f_i(x_i, x_{n(i)}), \quad (2)$$

and the minimum aggregate cost as $p^* = f(x^*)$

Dependence of individual cost function (1) on variables of the neighbors $x_{n(i)}$ requires that x_i and x_j for $j \in n(i)$ is chosen at the same time in order to achieve the minimum. For a connected network, cost minimization requires that all the variables are chosen simultaneously since an agent's variable is included on the functions of its neighbors and its neighbors' variables are also included on the functions of their neighbors and so on. However, simultaneous selection requires an unrealistic level of coordination among agents.

Instead, consider a model where each agent acts individually based on the information provided by its neighbors. Define a *rational* action for agent i as to select the value for its own variable that minimizes its local cost using available local information at time t,

$$\tilde{x}_i(t) = \operatorname*{argmin}_{x_i \in \mathcal{X}_i} f_i\big(x_i, x_{n(i)}(t)\big). \quad (3)$$

The update in (3) is based on local information and is executed solely by agent i, therefore it stands for a more realistic representation for network optimization. While possible, it is not always accurate to assume that agents follow rational policies perfectly. It is often the case that agents apply heuristic rules in their decision making which are prone to randomness and suboptimality. We introduce the concept of heuristic rational decision-making as a way to capture this type of sub-optimal behavior.

Definition 1 *We say that a probabilistic rule $x_i(t)$ is heuristic rational for agent i if and only if its expectation is a rational action defined as in* (3)*,*

$$\mathbb{E}\left[x_i(t)\right] = \tilde{x}_i(t) = \operatorname*{argmin}_{x_i \in \mathcal{X}_i} f_i\big(x_i, x_{n(i)}(t)\big) \quad (4)$$

We include a random activation rule that determines when each agent i modifies its variable according to a heuristic rational rule. For the activation rule, we use $k \in \mathbb{N}$ to index activations. The kth activation occurs at time t_k and involves a unique agent $i = i_k$ modifying variable $x_i = x_{i_k}$. Activations occur at random times and all agents have positive probability to become active in any given time interval. Variables $x_i(t_k)$ stay unchanged for all the other agents $i \neq i_k$. The variable $x_{i_k}(t_k)$ for terminal i_k is updated based on local observation of neighboring variables and follows heuristic rationality (4).

The following example demonstrates a possible application of heuristic rational network optimization on a field estimation problem with a wireless sensor network (WSN).

Example 1 (Field estimation using WSN) *Consider a network of sensors deployed on a spatially varying field. Each sensor i is responsible for estimating the value x_i that is associated with its location. This estimation is done via measurements y_i that are collected from sensor i's location. Measurements collected from the same location at different times are assumed to be independent given x_i and are corrupted with zero-mean Gaussian noise with variance σ^2, i.e. $\mathrm{P}\left(y_i \,\middle|\, x_i\right) = e^{-(y_i - x_i)^2/2\sigma^2}/\sqrt{2\pi\sigma^2}$. The network goal is specified to compute estimates that maximize log likelihood function $\ln P(x \,|\, y)$ where $y := \{y_i\}_{i \in V}$. Applying Bayes' rule, the estimates can be computed using the following formula*

$$\hat{x} = \operatorname{argmax} \ln \mathrm{P}\left(x \,\middle|\, y\right) := \operatorname{argmax} \big(\ln \mathrm{P}\left(y \,\middle|\, x\right) + \ln \mathrm{P}\left(x\right) \big). \quad (5)$$

Notice that the estimates in (5) *are coupled. In order to obtain a global objective function with decoupled structure similar to* (1)*, we consider Markov random fields (MRF) where spatial dependency is simplified to dependency on neighboring locations. Markovianity property affords that $\mathrm{P}\left(x_i \,\middle|\, x_j, j \neq i\right) = \mathrm{P}\left(x_i \,\middle|\, x_{n(i)}\right)$. According to Hammersley-Clifford theorem [9, Ch. 4.3], a MRF has an equivalent representation in terms of a Gibbs random field with the probability distribution that takes the form $\mathrm{P}\left(x\right) = \mathsf{e}^{-U(x)}/\alpha$ where α is some normalizing constant. $U(x)$ is called the field energy function and has the following structure $U(x) = \sum_{i, j \in n(i)} u_{ij}(x_i - x_j)$. When a Gaussian MRF is considered, field energy function becomes a quadratic function, $U(x) = \sum_{i, j \in n(i)} (x_i - x_j)^2/2\lambda$ with smoothing coefficient λ. Using this representation in* (5) *and further noting that since observations are conditionally independent $\mathrm{P}\left(y \,\middle|\, x\right) = \prod_i \mathrm{P}\left(y_i \,\middle|\, x_i\right)$ it follows that*

$$\hat{x} = \operatorname{argmin} \big(\sum_{i \in V} (y_i - x_i)^2/2\sigma^2 + \sum_{i \in V, j \in n(i)} (x_i - x_j)^2/2\lambda \big). \quad (6)$$

When we rearrange terms and define the functions $f_{0i}(\mathbf{x}_i) = (y_i - x_i)^2/2\sigma^2$ and $f_{ij}(x_i, x_j) = (x_i - x_j)^2/2\lambda$, minimand in (6) *has the same decoupled form as the global network cost in* (2) *and the local cost function becomes*

$$f_i(x_i, x_{n(i)}) = (y_i - x_i)^2/2\sigma^2 + \sum_{j \in n(i)} (x_i - x_j)^2/2\lambda. \quad (7)$$

A rational estimate for sensor i minimizing (7) *is to compute*

$$\tilde{x}_i = \frac{\lambda y_i + \sigma^2 \sum_{j \in n(i)} x_j}{\lambda + N_i \sigma^2} \quad (8)$$

given estimates of neighbors $x_{n(i)}$ and measurement y_i. Heuristic rationality can be used to account for communication errors, quantization effects during the local signal processing or the communication stages, and model mismatch; see Section V.

The goal in this paper is to characterize the performance of the sequence of iterates $x(t_k)$ generated by recursive application

of heuristic rational rules with respect to the optimal configuration x^*. For this goal, we define the stochastic process $\{F_k\}_{k\in\mathbb{N}}$ of optimality gaps with elements

$$F_k := f(x(t_k)) - p^*. \tag{9}$$

In Section III, we establish that the optimality gap F_k achieves a small value with probability 1 infinitely often (Theorem 1). Furthermore, it is shown that F_k stays close to this value with large probability (Section IV).

III. Near Optimality and Convergence

We analogize heuristic rational update rules with random activation scheme to a stochastic version of block coordinate descent on the function $f(x)$. Coordinate descent algorithms alternate between descents on different subsets of variables chosen according to a given rule [10]. We can identify agent's variables as coordinate blocks and random activation as the selection rule. The structure of the local cost function $f_i(x_i, x_{n(i)})$ in (1) is such that block coordinate descent can be applied in a distributed manner. Given the analogous decision model proposed, we show that indeed our expectations that heuristic rational iterates will achieve a cost close to aggregate optimal cost p^* is true. We need the following assumptions for our analysis.

(A1) The global cost $f(x)$ is strongly convex in that there exists a constant $m > 0$ such that for any pair of points $x \in \mathcal{X}$ and $y \in \mathcal{X}$ it holds

$$f(y) \geq f(x) + \nabla f(x)^T(y-x) + \frac{m}{2}||y-x||^2. \tag{10}$$

(A2) Gradients of the global cost $f(x)$ are Lipschitz in that there exists a constant $M > 0$ such that for any pair of points $x \in \mathcal{X}$ and $y \in \mathcal{X}$ it holds

$$f(y) \leq f(x) + \nabla f(x)^T(y-x) + \frac{M}{2}||y-x||^2. \tag{11}$$

(A3) At any given time t, all agents are equally likely to become active.

(A4) The mean square error of the heuristic rational action $x_{i_k}(t_k)$ with respect to the corresponding rational action $\tilde{x}_{i_k}(t_k)$ is bounded [cf. (4)].

$$\mathbb{E}\left[\|x_{i_k}(t_k) - \tilde{x}_{i_k}(t_k)\|^2\right] \leq \sigma^2. \tag{12}$$

The first two assumptions are typical in convergence analysis of descent algorithms. They are satisfied by Example 1 in Section II. An asynchronous update rule in which all agents having an activation clock based on independent exponential waiting times with equal means satisfies the third assumption. This assumption can easily be substituted by laxer conditions such as each agent having only a positive probability at any given time t and our results will continue to hold with minor modifications. Assumption (A4) caps the average irrationality of each agent by bounding the deviation from the rational decision (3). Note that this is a bound on mean square which means arbitrarily bad isolated actions are allowed. Furthermore, our results are parametric on the irrationality bound σ^2 that is the iterates of F_k are guaranteed to approach a close neighborhood of zero and closeness is proportional to σ^2.

Our first result shows a contraction for the optimality gap F_{k+1} at time t_{k+1} in reference to the optimality gap F_k at time t_k – see [11] for proof.

Lemma 1 *Consider iterates of heuristic rational rule $x(t_k) = \{x_i(t_k)\}_{i\in V}$ such that at time t_k agent i_k updates her local variables $x_i(t_k)$. The optimality gaps F_{k+1} and F_k as defined in* (9) *satisfy*

$$\mathbb{E}\left[F_{k+1} \,\middle|\, \mathbf{x}(k)\right] - F_k \leq -\beta F_k + \frac{M\sigma^2}{2} \tag{13}$$

where we defined the condition number $\beta := m/(MN)$ and used the shorthand notation $x(k) := x(t_k)$.

The inequality (13) indicates that F_k is behaves like a supermartingale when the right hand side is less than zero. This happens when F_k is far away or more precisely when $F_k \geq M\sigma^2/2\beta$. In this case, the rationality dominates the aggregate behavior and we have the expectation to descent toward optimality. This expected rationality prevails over a larger range of values of F_k when σ^2 is small. When the process is within the boundary i.e. $F_k < M\sigma^2/2\beta$, the inequality cannot not claim an expected decrease. This case can be interpreted as operating at the region where irrationality prevails.

Considering that the process is attracted towards zero, when optimality gap is larger than $M\sigma^2/2\beta$, we expect to see the process F_k becoming smaller than $M\sigma^2/2\beta$ at least once. This intuition is correct as we state in the following theorem – see [11] for proof.

Theorem 1 *Let $F_k^{\text{best}} := \min_{l\in[0,k]} F_l$ be the best optimality gap by time t_k. If assumptions (A1)-(A4) hold, it follows that*

$$\lim_{k\to\infty} F_k^{\text{best}} \leq \frac{M\sigma^2}{2\beta}. \quad \textit{a.s.} \tag{14}$$

Theorem 1 confirms that the optimality gap becomes smaller than $M\sigma^2/2\beta$ at least once for almost all realizations. It also implies that this happens infinitely often which means that the process moves outside the neighborhood of optimality defined by $M\sigma^2/2\beta$ infinitely often. Hence, it is worthwhile to try to characterize process' behavior when it is outside the near optimality region.

IV. Excursions outside near Optimality

Theorem 1 provides the motivation to look at excursions of the optimality gap when it is outside the near optimality region because it has no control over the process inside the near optimality region. Inside the near optimality region, irrationality dominates which guarantees to throw optimality gap outside near optimality. When the process is outside near optimality, the only characterization that is useful is that it behaves like a supermartingale shown in Lemma 1. The inequality (13) in Lemma 1 indicates that the expected descent increases as the optimality gap grows. This gives us the leverage for showing that the excursions from near optimality are bounded exponentially.

Next, we formally define excursions away from the optimality neighborhood. Suppose that at given iteration k, the optimality

gap is $F_k = (1+\rho)M\sigma^2/2\beta$, i.e., larger than the neighborhood border by a factor $\rho > 0$. Further consider a given value $\gamma > F_k$. We define excursion as the trajectory $F_k, F_{k+1}, \ldots, F_{k+L}$ of the optimality gap until the process returns to a value $F_{k+L} < F_k$ smaller than the given gap F_k from which the excursion started. Notice that L is a random stopping time given by the first time the process achieves a smaller gap than the starting point $L = \min_{l>0}\big(F_{k+l} < F_k\big)$. In particular, we are interested in the worst value $F_k^\dagger = \max(F_k, F_{k+1}, \ldots, F_{k+L})$ reached during the excursion.

Our goal here is to determine the probability that the worst value attained during the excursion exceeds the given γ, $\mathrm{P}\big(F_k^\dagger \geq \gamma\big)$. We need the following additional assumption to obtain bound on the probability $\mathrm{P}\big(F_k^\dagger \geq \gamma\big)$.

(A5) The difference on optimality gaps between successive iterations is almost surely bounded by a finite constant $\kappa > 0$, i.e., for all times k we have that

$$\mathrm{P}\left(|F_{k+1} - F_k| \leq \kappa \,\middle|\, F_k\right) = 1. \tag{15}$$

Assumption (A5) is satisfied when the functions $f_{ij}(x_i, x_j)$ are bounded for all feasible values $x_i \in \mathcal{X}_i$ and $x_j \in \mathcal{X}_j$. Assumption (A5) can be satisfied if the differences $\|x_{i_k}(t_k) - \tilde{x}_{i_k}(t_k)\|$ between rational and heuristic rational actions are almost surely bounded. Note that in this case, Assumption (A4) is automatically satisfied. For the Example 1 in Section II, Assumption (A5) is satisfied if the range of values that can be measured by sensors is bounded. In the following theorem, we state the exponential bound on $\mathrm{P}\big(F_k^\dagger \geq \gamma\big)$ – see [11] for proof.

Theorem 2 *Assume that at time k_0 the excursion starts at a value that is away from the near optimality boundary established in Theorem 1 by a factor of $\rho > 0$, i.e., $F_{k_0} = (1+\rho)M\sigma^2/2\beta$. If assumptions (A1)-(A5) hold, then, for arbitrary given constant γ, we have*

$$\mathrm{P}\left(F_{k_0}^\dagger \geq \gamma \middle| F_{k_0}\right) \leq \mathrm{e}^{-c(\gamma - F_{k_0})}, \tag{16}$$

with $c = 2\rho M\sigma^2/[(\rho M\sigma^2)^2 + \kappa^2]$.

Theorem 2 indicates that the probability of worst optimality gap during an excursion decreases exponentially as γ increases. This shows that being arbitrarily far away from near optimality is highly unlikely as it is a bound on worst optimality gap reached during the excursion process. The exponential bound in (16) is also dependent on a coefficient c that scales the exponential term. As the coefficient c increases, the bound on the right hand side of (16) decreases making it stronger. Coefficient c depends on constants such as mean square error bound σ^2 and optimality gap increment bound κ. We observe that increase in any one of these constants has an inverse effect on the scaling coefficient which means a laxer bound on the right hand side.

Next, we give numerical examples for an estimation problem of a MRF using WSN in which sensors follow a heuristic rational rule.

V. Simulation

Consider a temperature field that has the spatial dependency structure of Gaussian MRF as explained in Example 1 of Section II. Further, there exists a network of temperature sensors placed on this field obtaining noisy temperature measurements from various locations.

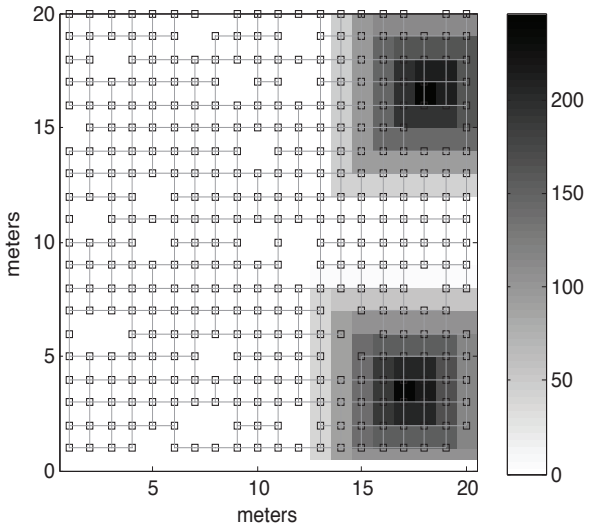

Fig. 1. Estimation of a temperature field with a sensor network. Two heat sources located at $(17\mathrm{m}, 4\mathrm{m})$ and $(18\mathrm{m}, 17\mathrm{m})$ set temperatures at their locations to $255^\circ F$. Temperature around these sources falls at a rate of $50^\circ F/\mathrm{m}$ and is set to $0^\circ F$ at other locations. The resulting temperature field is encoded according to the scale on the right. A sensor network with $N = 370$ is deployed to estimate this field with lines representing network connectivity. Sensors estimate the temperature at their location.

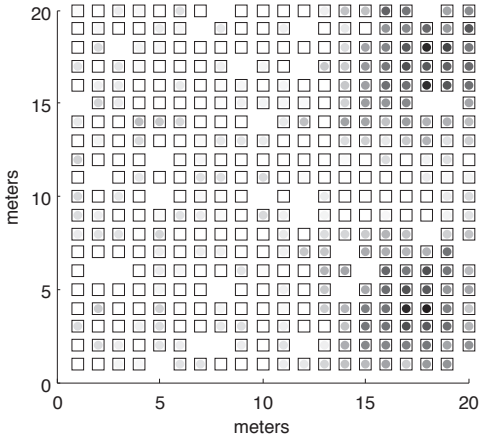

Fig. 2. Sensor temperature field estimates at time $t = 0$ based only on initial measurements. Gaussian MRF has smoothing coefficient $\lambda = 10^3$. Observation noise power is equal to $\sigma^2 = 0.9 \times 10^3$. Sensors update their temperature estimates using the heuristic rational update in (17) with quantization level $\Delta = 1$. Estimates are encoded using the same scale of Fig. 1.

Due to bandwidth limitations sensors quantize their estimations before transmitting them to their neighbors. Thus, the signal received by sensor i from sensor j is a quantized version $x_{qj}(t)$ that can be written as $x_{qj}(t) = x_j(t) + q_j(t)$ for some quantization noise $q_j(t)$. The update in (8) is therefore not the one carried out by sensor i. Rather, the estimate of sensor i is

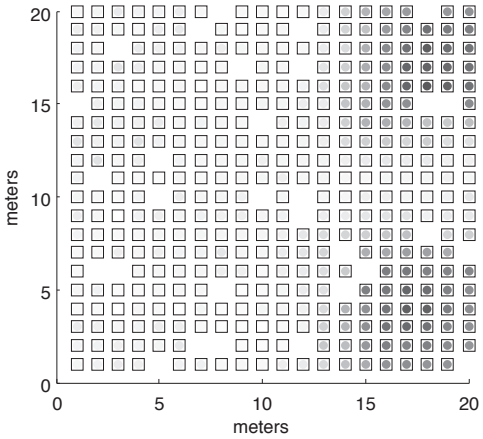

Fig. 3. Sensor temperature field estimates at time $t = 4$. As time progresses sensor estimates become closer to field values as information from neighboring nodes is incorporated into local estimates.

updated to

$$x_i(t) = \frac{\lambda y_i + \sigma^2 \sum_{j \in n(i)} x_{qj}(t)}{\lambda + N_i \sigma^2} \ . \tag{17}$$

Assuming the quantization noise has zero mean, i.e., $\mathbb{E}[q_i(t)] = 0$, it follows that $\mathbb{E}[x_i(t)] = \tilde{x}_i(t)$. We can then think of the update in (17) as a heuristic rational version of the update in (8).

Numerical simulations for a WSN with $N = 370$ temperature sensors performing the update in (17) to estimate a temperature field are shown in figs. 1 - 4. The temperature field is generated as a 20m$\times$20m grid with temperature values ranging in $[0°F, 255°F]$. Two heat sources at locations $\mathbf{h}_1 = (17, 4)$ and $\mathbf{h}_2 = (18, 17)$ set the temperature in these points to $255°F$. Heat dissipates from these sources for a range of $5m$ at a rate of $50°F$/m as one moves away from the heat source. The sources do not influence the temperature outside $5m$ range and the temperature is set to $0°F$ for those locations; see Fig. 1. The sensors are located at random positions in a 1m grid with communication between sensors occurring only between sensors located less than 1m apart; see Fig. 1. The temperature field is a uniform Gaussian MRF with smoothing parameter $\lambda = 10^3$ and observation noise as Gaussian with variance $\sigma^2 = 0.9 \times 10^3$. The quantization levels for temperature estimates are integers in $[0, 255]$. Each sensor activates according to an exponential distribution with mean one. Hence, the expected number of updates by time $t = 1$ is equal to the number of sensors N.

Figs. 2-3 use squares to label locations on which there is a sensor. The estimates follow color encoding in Fig. 1. Figs. 2-3 display sensor estimates at times $t = 0$, and $t = 4$, respectively. At time $t = 0$ estimates are based on local observations only and thereby can show significant difference with respect to field values. By time $t = 4$, we observe that the sensor estimates are refined and are closer to original value as shown in Fig. 1. We also observe that sensor estimates close to heat sources are affected by their neighboring sensors that are located on lower temperature points. Furthermore, sensors that lie in between two heat regions (between points (14,9) and (20,12)) tend to have larger estimates than the correct value of zero. These are all caused by the mismatch between the underlying modeling assumption of correlated field and what the field actually entails.

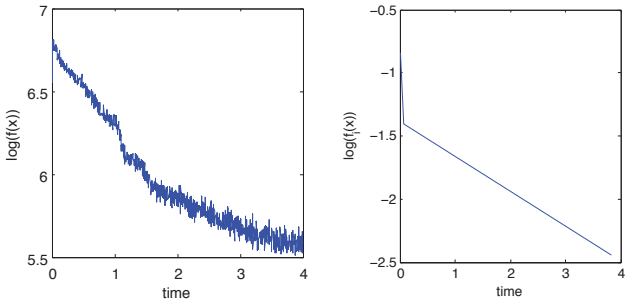

(a) (b)

Fig. 4. Global (a) and individual (b) log-likelihoods for temperature field estimation.

Figs. 4(a) and 4(b) show the evolution of the global log-likelihood function $f(\mathbf{x}) = \sum_i f_i(x_i(t), \mathbf{x}_{n(i)}(t))$ and the local cost function $f_i(x_i(t), \mathbf{x}_{n(i)}(t))$ in (7) for a selected i. The log-likelihoods tend to decrease thereby resulting on more refined estimates.

VI. Conclusion

We considered a distributed approach to a network optimization problem where agents act according to a heuristic rational rule that is on the average optimal. For this setup, we show that the minimum global network cost value converges to near optimality almost surely. Furthermore, we show that the excursions from near optimality is exponentially bounded.

References

[1] F. Cattivelli and A.H. Sayed. Modeling bird flight formations using diffusion adaptation. *IEEE Trans. on Signal Processing*, 59(5), 2011.

[2] V.D. Blondel, J.M. Hendrickx, A. Olshevsky, and Tsitsiklis J.N. Convergence in multiagent coordination, consensus, and flocking. In *Proc. of IEEE CDC*, 2005.

[3] A. Jadbabaie, J. Lin, and A.S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Trans. on Automatic Control*, 48(6):988–1001, 2003.

[4] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Gossip algorithms: Design, analysis, and applications. In *Proceedings of IEEE INFOCOM*, pages 237–252, 1998.

[5] E. Mossel and G. Schonenebeck. Reaching consensus on social networks. In *Proc. of the ICS '10*, 2010.

[6] A. Nedic, A. Ozdaglar, and A.P. Parrilo. Constrained consensus and optimization in multi-agent networks. *IEEE Transactions on Automatic Control*, 55(4):922–938, 2010.

[7] M. Rabbat, R. Nowak, and J. Bucklew. Generalized consensus computation in networked systems with erasure links. In *IEEE Workshop on SPAWC Proc.*, June 2005.

[8] S.S. Ram, A. Nedic, and Veeravalli V.V. Distributed stochastic subgradient projection algorithms for convex optimization. *Journal of Optimization Theory and Applications*, 147(3):516–545, 2010.

[9] N.A.C. Cressie. *Statistics for Spatial Data*. New York: Wiley, revise edition, 1993.

[10] P. Tseng. Convergence of a block coordinate descent method for non-differentiable minimization. *Convergence of a Block Coordinate Descent Method for Nondifferentiable Minimization*, 109(3):475–494, 2001.

[11] C. Eksin and A. Ribeiro. Distributed network optimization with heuristic rational agents. *IEEE Trans. Signal Process*, (submitted) [Online]. Available: http://www.seas.upenn.edu/~aribeiro/preprints/hra_main.pdf, 2011.