

Deep Q-Learning for Self-Organizing Networks Fault Management and Radio Performance Improvement

Faris B. Mismar and Brian L. Evans

Wireless Networking and Communications Group, The University of Texas at Austin, Austin, TX 78712 USA

Abstract—We propose an algorithm to automate fault management in an outdoor cellular network using deep reinforcement learning (RL) against wireless impairments. This algorithm enables the cellular network cluster to self-heal by allowing RL to learn how to improve the downlink signal to interference plus noise ratio through exploration and exploitation of various alarm corrective actions. The main contributions of this paper are to 1) introduce a deep RL-based fault handling algorithm which self-organizing networks can implement in a polynomial runtime and 2) show that this fault management method can improve the radio link performance in a realistic network setup. Simulation results show that our proposed algorithm learns an action sequence to clear alarms and improve the performance in the cellular cluster better than existing algorithms, even against the randomness of the network fault occurrences and user movements.

Index Terms—reinforcement learning, wireless, tuning, optimization, artificial intelligence, SON.

I. INTRODUCTION

Self-organizing networks (SON) are expected to improve the efficiency of cellular coverage tuning to meet target service performance metrics [1]. Industry standards [2] refer to the category of SON which performs automatic handling of network faults as *self-healing*. Management of network faults is one of the functional areas defined in an industry framework [3] known as FCAPS: fault, configuration, accounting, performance, and security. Fault management detects faults and corrects them. Wireless network faults can be corrected through parameter adjustments which optimize the network performance. This enables it to carry the traffic it has been dimensioned for with high reliability and end-user *quality of experience* (QoE). However, this is also a perpetual and costly task.

We use deep *reinforcement learning* (RL) where the SON learns fault management with no human supervision. Our proposed addition of deep RL to the SON is shown in Fig. 1. SON aided with the deep RL-based algorithm monitors the performance data of an outdoor cellular network and analyzes it to derive proper parameter corrective steps and implements them. For this purpose, we model a cellular cluster of base stations in an outdoor environment with mobile devices scattered in the vicinity with the focus on the downlink. We refer to any of these mobile devices as a *user equipment* (UE). We chose the *signal to interference and noise ratio* (SINR) and throughput as measures of cellular capacity. We focus on the *fourth generation of wireless communications* or *Long Term Evolution* (4G LTE) and

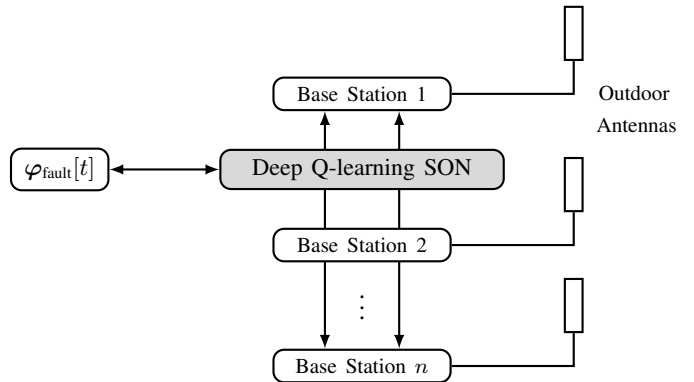


Fig. 1. The Deep Q-learning self-organizing network (SON) module interacting with several outdoor base stations.

its variants since the system-level simulator [4] supports them. The results can be applied to any similar *orthogonal frequency-division multiplexing* (OFDM) including the *fifth generation of wireless communications new radio* (5G NR) where SON has a highly anticipated role [1].

The first deep RL framework to learn control policies using reinforcement learning was introduced in [5]. This framework outperformed human experts. The authors in [6] used Q-learning as part of their SON implementation for mobile load balancing and mobility optimization for cell reselection and handovers. They used cells with single antennas. We on the other hand use *multiple-input multiple-output* (MIMO) in our transceivers—a fundamental setup for present and futuristic network deployments.

Deep learning in mobile and wireless networking with interference alignment was studied in [7]. Relaxed *channel state information* (CSI) assumptions were made in the study where the CSI transition matrix was identical across all users. We do not make this assumption since we focus on upper layers in the wireless stack. A two-dimensional convolutional neural network was used in simulations, which imposes unfounded spatially invariant relationships between learning features, or local connection patterns [8]. We avoid this in our design of our deep neural network. The authors in [9] provided a means to improve the handover execution success rate using supervised machine learning but did not use reinforcement learning which has the ability to learn from previous actions. In [10], a method was proposed to extract




Fig. 2. Structure of the neural network used for the DQN implementation with two hidden layers each of dimension H . $|\mathcal{S}| = m$, $|\mathcal{A}| = n$.

the knowledge base from solved fault troubleshooting cases using data mining and supervised learning techniques using fuzzy logic. Expert opinion was used to define performance measurements and target values. We on the other hand use RL to derive a policy to map actions to be taken by the self-healing functionality in response to select common number of faults in the network. The method in [11] showed that deep RL can be run in a distributed fashion.

Our main contributions are as follows:

- Introduce a deep RL-based fault handling algorithm which self-organizing networks can implement in a polynomial runtime.
- Show that this fault management method can improve the radio link performance in a realistic outdoor network setup.

II. SYSTEM MODEL

The system comprises a wireless network of a macro base station operating in the sub-6 GHz frequency range with a single tier of surrounding macro base stations, and a machine learning algorithm using deep RL which could reside in the serving base station or at a central location.

A. Network

The network is an outdoor cellular cluster using frequency division duplex and multiple access with one tier of neighboring cells each with a hexagonal geometry and an inter-site distance of length L . All of these cells are neighbors to one another. The UEs are equipped with one antenna each and are allowed to perform handovers between cells. The cells have multiple antennas each. Any UE is served by a serving cell c , which is a member of the set of cells \mathcal{C} . In this network, there are $q|\mathcal{C}|$ UEs, where q is the number of UEs per cell.

This cellular network can be in a normal operational state or undergo several issues or faults. The set of faults are $\mathcal{N} = \{\nu_i\}_{i=0}^{|\mathcal{N}|-1}$. We choose a few common faults which can be resolved by SON and set their rate of occurrence as in Table I.

B. Reinforcement Learning

We use Q -learning with a deep neural network as in [5]. This deep neural network is known as the *deep Q-network* (DQN) and is shown in Fig. 2. We use DQN with




Fig. 3. Reinforcement learning elements. FM stands for fault management.




Fig. 4. Markov decision process used in the formulation of the SON fault handling algorithm.

experience replay and formulate the SON fault management as a deep RL problem as in Algorithm 1. The framing of the reinforcement learning problem is shown in Fig. 3.

To formulate this problem as a reinforcement learning problem, we define a *Markov decision process* (MDP) with transition probabilities $p(\cdot)$ as in Fig. 4. The set of actions carried out by the self-healing agent are $\mathcal{A} = \{a_i\}_{i=0}^{n-1}$ and the set of network states are $\mathcal{S} = \{s_i\}_{i=0}^{m-1}$. These are shown in Table II. A state s is *terminal* if it is the final state or if the objective has been met. We define the reward as:

$$r_{s,a}[t] \triangleq \begin{cases} r_1, & 0 < |\varphi_{\text{fault}}[t-1]| < |\varphi_{\text{fault}}[t]| \\ r_2, & |\varphi_{\text{fault}}[t]| = |\varphi_{\text{fault}}[t-1]| \\ r_3, & |\varphi_{\text{fault}}[t]| < |\varphi_{\text{fault}}[t-1]| \\ r_4, & |\varphi_{\text{fault}}[t]| = 0 \text{ (i.e., objective is met)} \end{cases} \quad (1)$$

where $|\varphi_{\text{fault}}[t]|$ is the number of bits that are set to logic-1 in the fault register at a given TTI t . The fault register $\varphi_{\text{fault}}[t]$ is a register of u boolean entries, where the i -th entry in the register corresponds to the fault with identifier i triggered in this cluster. It is initialized to all logic-0 and set whenever a fault i happens in the network and unset only when all similar faults are cleared. These faults are shown in Table I.

With a network having $|\mathcal{C}|$ cells, n alarm-clearing actions, and m states, the number of elements in a table required is in $\mathcal{O}(mn|\mathcal{C}|)$. In networks with thousands of cells and alarms, the tabular Q -learning method to keep track of the state-action values in a table does not scale due to this multiplicative nature.

The *episode* z : $z \in \{0, 1, \dots, \zeta\}$ is a period of time for an agent-environment interaction to take place. During an episode, the deep RL agent gathers *experience* from a finite number of episodes where each episode has a duration of τ *transmit time intervals* (TTIs). This agent stores the experience $e[t]$ defined as $(s, a, r_{s,a}, s')$ in a dataset called the *replay memory* \mathcal{D} [5].

We next define the estimated function $Q^*(s, a)$, which is the expected discounted reward when starting in state s and selecting an action a as [5]:

TABLE I
NETWORK FAULTS \mathcal{N}

Action ν	Definition	Rate	Action ν	Definition	Rate
0	Cluster is normal.	p_0	5	Feeder fault alarm cleared. [†]	p_5
1	Changed antenna azimuth clockwise (e.g., due to wind).	p_1	6	Neighboring cell is up again. [†]	p_6
2	Neighboring cell is down.	p_2	7	Transmit diversity is normal. [†]	p_7
3	Transmit diversity failed.	p_3	8	Reset antenna azimuth. [†]	p_8
4	Feeder fault alarm (3 dB loss of signal).	p_4			

[†] These actions cannot happen if their respective alarm did not happen first.

$$Q^*(s, a) \triangleq \mathbb{E}_{s'} \left[r_{s,a} + \gamma \max_{a'} Q^*(s', a') \mid s, a \right] \quad (2)$$

If we define a DQN with its weights at time t as θ_t , then (2) can be approximated as $Q(s, a; \theta_t) \approx Q^*(s, a)$ as $t \rightarrow \infty$. This DQN is trained through minimizing the *mean squared error* (MSE) convex loss function

$$L_t(\theta_t) \triangleq \mathbb{E}_{s,a} [(y_t - Q(s, a; \theta_t))^2] \quad (3)$$

where y_t is an estimate obtained from the DQN using its weights at time $t - 1$ as:

$$y_t \triangleq \mathbb{E}_{s'} \left[r_{s,a} + \gamma \max_{a'} Q(s', a'; \theta_{t-1}) \mid s, a \right]. \quad (4)$$

The weights θ are updated after every iteration in time using a method of the stochastic gradient descent algorithm (SGD) called “adaptive moments” [13]. We also use the *rectified linear unit* (ReLU) $x \mapsto \max(0, x)$ as the activation function of each node in the DQN. This deep learning process repeats for all the episodes z .

During an episode z , the deep RL agent tries to maximize the total value of the discounted rewards it receives in response to its action. It uses a near-greedy action selection rule. This is because with large number of episodes ζ , every actions will have been sampled enough for the state-action value function to converge [12]. The near-greedy action selection rule of choice is the ϵ -greedy strategy for learning. In this strategy, ϵ is the *exploration rate* and serves to select a random action $a \in \mathcal{A}$ with a probability ϵ : $0 < \epsilon < 1$, known as exploration, or selects an action from the replay memory \mathcal{D} , which is also known as exploitation, with a probability $1 - \epsilon$. The exploration rate decays in every episode until it reaches the minimum exploration rate ϵ_{\min} .

III. FAULT HANDLING ALGORITHMS

A. Random

The faults are cleared from a random sample of active faults. We choose the discrete uniform random distribution for the healing of the faults in the network since the discrete uniform distribution maximizes the discrete entropy [14]. This is a non-trivial lower bound of performance compared to taking no action.

B. First-In, First-Out

In this approach, the SON takes fault handling actions immediately in the next TTI in the order these faults happen.

Algorithm 1: SON Fault Management

Input: The set of fault handling actions \mathcal{A} in a network \mathcal{C} .

Output: Optimal fault handling commands given during a frame z , which has a duration of τ .

- 1 Define the fault management states \mathcal{S} , the exploration rate ϵ , the decay rate d , the discount factor γ , and minimum exploration rate ϵ_{\min} .
- 2 Initialize time, states, actions, fault handling register, and replay memory \mathcal{D} .
- 3 **repeat**
- 4 $t := t + 1$
- 5 $\epsilon := \max(\epsilon \cdot d, \epsilon_{\min})$
- 6 Sample $r \sim \text{Uniform}(0, 1)$
- 7 **if** $r \leq \epsilon$ **then**
- 8 Select an action $a \in \mathcal{A}$ at random.
- 9 **else**
- 10 Select an action $a = \arg \max_{a'} Q(s, a'; \theta_t)$.
- 11 **end**
- 12 Perform action a to resolve alarm and update $\varphi_{\text{fault}}[t]$.
- 13 Obtain reward $r_{s,a}$ from (1).
- 14 Observe next state s' .
- 15 Store experience $e[t] \triangleq (s, a, r_{s,a}, s')$ in \mathcal{D} .
- 16 Sample from \mathcal{D} for experience $e_j \triangleq (s_j, a_j, r_j, s_{j+1})$.
- 17 **if** s_{j+1} is terminal **then**
- 18 Set $y_j := r_j$
- 19 **else**
- 20 Set $y_j := r_j + \gamma \max_{a'} Q(s_{j+1}, a'; \theta_t)$
- 21 **end**
- 22 Perform SGD on $(y_j - Q(s_j, a_j; \theta_t))^2$
- 23 $s := s'$
- 24 **until** $|\varphi_{\text{fault}}[t]| = 0$ or $t \geq \tau$
- 25 Proceed to the next LTE-A frame $z + 1$.

C. Proposed

Our proposed deep RL-based algorithm is shown in Algorithm 1. Unlike FIFO, this algorithm can handle faults regardless of the order they happen due to the ϵ -greedy learning strategy.

We next compute the time complexities of these various fault handling algorithms:

- For the random algorithm, an action is randomly sampled from a list of actions, therefore it has a time complexity in $\mathcal{O}(1)$ per iteration or $\mathcal{O}(\tau)$ total.
- The First-In First-Out (FIFO) fault-handling algorithm reviews the alarm register every TTI and therefore has a time complexity in $\mathcal{O}(\max(u, |\mathcal{C}|))$ per iteration or $\mathcal{O}(\tau \max(u, |\mathcal{C}|))$ total.
- For our proposed algorithm, the time complexity of the for DQN backpropagation algorithm is at least in

TABLE II
SON FAULT MANAGEMENT ALGORITHM – SIMULATED ACTIONS \mathcal{A} AND STATES \mathcal{S}

Action a	Definition	State s	Definition
0	No actions issued.	0	No actions issued (a transient state).
1	Faulty neighbor cell is up again.	1	Number of active alarms has increased.
2	Serving cell transmit diversity enabled.	2	Number of active alarms has decreased.
3	Serving cell losses recovered.		
4	Serving cell azimuth set to default value.		

TABLE III
MACHINE LEARNING PARAMETERS

Parameter	Value
Neural network hidden layer width H	24
Loss function $L_t(\theta_t)$	mean squared error
Optimizer	[13]
Number of episodes ζ	50
Discount factor γ	0.950
Exploration rate ϵ	1.000
Exploration rate decay d	0.91
Minimum exploration rate ϵ_{\min}	0.010
One episode duration τ	20 TTIs
Number of states	3
Number of actions	5

$\mathcal{O}(k(\theta)\tau\zeta|\mathcal{C}||\mathcal{A}|)$ [15], where $k(\theta)$ is a function of the depth and number of the hidden layers θ .

Although our proposed algorithm has the highest time complexity cost, the complexity is not dependent on the number of UEs being served, and therefore it is scalable in the number of UEs served in a cluster. Furthermore, while the random algorithm can be trivially distributed across multiple cells independently and require a space complexity in $\mathcal{O}(|\mathcal{A}|)$, the distributed implementation is less scalable for the FIFO model since it requires $|\mathcal{C}|(|\mathcal{C}| - 1)/2$ links to communicate about faults and fault management leaving us with a message passing complexity in $\mathcal{O}(|\mathcal{C}|^2)$. Our proposed algorithm can be run in a distributed fashion owed to its deep learning component [11] and centralized collection point at the SON leaving us with message passing complexity in $\mathcal{O}(|\mathcal{C}|)$. We refer to the source code [16] for further implementation details.

IV. PERFORMANCE METRICS

We evaluate the downlink SINR with power allocation using the waterfilling algorithm at the transmitter and the zero-forcing equalization at the receiver as in [4]. We also use the average downlink cell throughput and the downlink user throughput, which is derived from its empirical cumulative distribution function (CDF) as follows: peak (95%), average, and edge (5%) [17].

V. SIMULATION RESULTS

We use the Vienna LTE-A Downlink System Level Simulator to simulate the outdoor network in Fig. 5 with reproducibility [4]. The choice between LTE-A or 5G in the sub-6 GHz range is driven by the OFDM numerology scaling factor. We implement Algorithm 1 using both MATLAB and Python [16]. The parameters used in our simulation are summarized in Tables III and IV. Each episode has a

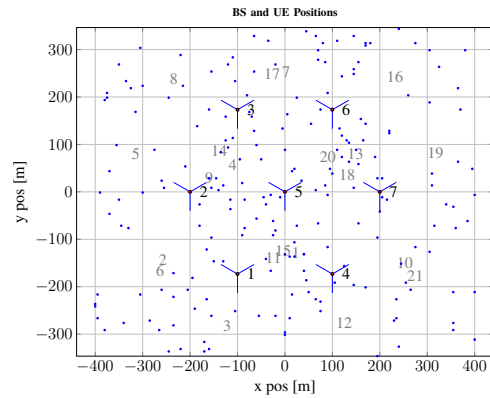


Fig. 5. The simulated outdoor network.

duration of $\tau = 50$ TTIs. In LTE-A or 5G, the duration of 1 TTI is equal to 1 ms. We set the rate of occurrences of faults in Table I as $p_0 = 5/9, p_1 = p_2 = p_3 = p_4 = 1/9$. This way we give all faults an equally likely chance of occurrence which is the worst case. For the rewards in (1), we set $r_1 = -1, r_2 = 0, r_3 = 1, r_4 = 5$.

In Fig. 6, we show the average downlink SINR distribution for all three algorithms for $q = 10$. While the FIFO algorithm performs incrementally better than the random algorithm, our proposed algorithm has a slightly higher downlink SINR overall and a significantly higher SINR gain in the range from 2 to 7 dB. The random algorithm leads to the worst performance due to the poor handling of faults. For our proposed algorithm, we show the results after $\zeta = 50$ episodes, which is after a total of $\tau\zeta = 1$ second.

The quantitative results of our simulation are shown in Table V. We observe that when the cell serves a low number of users (i.e., low cell load), our proposed algorithm outperforms the random algorithm as a lower bound and outperforms the FIFO algorithm. However, as the cell serves more UEs ($q = 50$), the performance of all algorithms becomes similar since the cellular resources are near depleted (i.e., a bottleneck) at highload and therefore clearing alarms does not lead to significant improvements.

VI. CONCLUSION

In this paper, we attempted to find an optimal solution policy for a RL based fault management algorithm in an outdoor realistic cellular environment. We motivated the need for deep RL in resolving faults in similar cellular environments. The use of deep RL allows the distributed implementation of the fault handling algorithm, which is

TABLE IV
RADIO ENVIRONMENT PARAMETERS

Parameter	Value	Parameter	Value
Bandwidth	10 MHz	Downlink center frequency	2100 MHz
LTE cyclic prefix	Normal	Cellular geometry	Hexagonal
Inter-site distance	200m	Scheduling algorithm	Proportional Fair
Number of cells in the network	21	Propagation model	COST231
Propagation environment	Urban	Number of concurrently active UEs per cell q^\dagger	{10, 50}
BS antenna model [†]	3gpp 36.942	BS maximum transmit power	46 dBm
BS antenna height	25 m	BS antenna electrical tilt	4°
BS MIMO configuration (# Tx, # Rx antennas)	(4, 2)	Noise power density	-174 dBm/Hz
UE average movement speed	3 km/h	UE height	1.5 m
Shadow fading margin standard deviation	8 dB	BS number of sectors per site	3

[†] BS is short for *base station* and UE is short for *user equipment*.

TABLE V
THROUGHPUT PERFORMANCE PER ALGORITHM FOR DIFFERENT NUMBER OF USERS PER CELL
User Equipment (UE) Throughput [Mbps]

UEs per cell	Random			FIFO			Proposed			Cell Throughput [Mbps]		
	Peak	Average	Edge	Peak	Average	Edge	Peak	Average	Edge	Average		
$q = 10$	3.48	1.78	0.53	3.52	1.79	0.54	3.55	1.84	0.58	17.77	17.95	18.37
$q = 50$	0.68	0.38	0.13	0.68	0.38	0.13	0.68	0.38	0.13	18.81	18.96	18.97




Fig. 6. The average downlink SINR empirical CDF as measured by the user equipment (UEs) for all three algorithms ($q = 10$).

useful when the cluster size is large. Our proposed solution works by allowing RL to learn how to improve the DL SINR through exploration and exploitation of various alarm corrective actions. The simulations showed that the deep RL-based method, which we proposed, can improve the performance of the cellular network measured by the downlink SINR and downlink throughputs. Therefore, the proposed deep RL-based automated cellular network tuning framework is useful in maintaining the end-user QoE in a network with impairments and faults.

REFERENCES

- [1] A. Imran, A. Zoha, and A. Abu-Dayya, "Challenges in 5G: how to empower SON with big data for enabling 5G," *IEEE Net.*, Nov. 2014.
- [2] 3GPP, "Telecommunication Management; Self-Organizing Networks (SON); Concepts and requirements," 3rd Generation Partnership
- [3] ITU-T, *Principles for a telecommunications management network*, International Telecommunication Union Std. M.3010, Feb. 2000. [Online]. Available: <https://www.itu.int/rec/T-REC-M.3010-200002-I/en>
- [4] M. Rupp, S. Schwarz, and M. Tarantetz, *The Vienna LTE-Advanced Simulators: Up and Downlink, Link and System Level Simulation*, 1st ed., ser. Signals and Communication Technology, 2016.
- [5] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing Atari with Deep Reinforcement Learning," ser. NIPS Deep Learning Workshop, 2013.
- [6] S. S. Mwanje, L. C. Schmelz, and A. Mitschele-Thiel, "Cognitive Cellular Networks: A Q-Learning Framework for Self-Organizing Networks," *IEEE Trans. Netw. Service Manag.*, Mar. 2016.
- [7] Y. He, Z. Zhang, F. R. Yu, N. Zhao, H. Yin, V. C. M. Leung, and Y. Zhang, "Deep-Reinforcement-Learning-Based Optimization for Cache-Enabled Opportunistic Interference Alignment Wireless Networks," *IEEE Transactions on Vehicular Technology*, Sep. 2017.
- [8] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. of the IEEE*, 1998.
- [9] F. B. Mismar and B. L. Evans, "Partially Blind Handovers for mmWave New Radio Aided by Sub-6 GHz LTE Signaling," in *IEEE Int. Conf. on Comm. Workshop on Evolutional Tech. & Ecosystems for 5G Phase II*, Jul. 2018.
- [10] E. J. Khatib, R. Barco, A. Gómez-Andrades, P. Muñoz, and I. Serrano, "Data mining for fuzzy diagnosis systems in LTE networks," *Expert Systems with Applications*, no. 21, pp. 7549–7559, Nov. 2015.
- [11] J. Dean, G. S. Corrado, R. Monga, K. Chen, M. Devin, Q. V. Le, M. Z. Mao, M. Ranzato, A. Senior, P. Tucker, K. Yang, and A. Y. Ng, "Large scale distributed deep networks," in *Proc. Int. Conf. on Neural Information Processing Systems*, ser. NIPS, 2012.
- [12] R. S. Sutton and A. G. Barto, *Introduction to Reinforcement Learning*, 1st ed. Cambridge, MA, USA: MIT Press, 1998.
- [13] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *Proc. Int. Conf. on Learning Representations*, 2015.
- [14] T. Cover and J. Thomas, *Elements of Information Theory*, 2006.
- [15] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion et al., "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, 2011.
- [16] F. B. Mismar. Deep Q-Learning SON Performance Code. [Online]. Available: <https://github.com/farismismar/Deep-Q-Learning-SON-Perf-Improvement>
- [17] Vienna Simulators LTE-A Downlink System Simulator. [Online]. Available: https://www.nt.tuwien.ac.at/wp-content/uploads/2015/11/LTEsystemDoc_v1_9Q2_2016.pdf