

Controlling the Movement of Robotic Bodyguards for Maximal Physical Protection

Taranjeet Singh Bhatia, Gürkan Solmaz, Damla Turgut and Ladislau Bölöni

Department of Electrical Engineering and Computer Science

University of Central Florida

4000 Central Florida Blvd, Orlando FL 32816

{tsbhatia, gsolmaz, turgut, lboloni}@eecs.ucf.edu

Abstract

In this paper we focus on a scenario where one or several robotic bodyguards protect a VIP moving in a public environment from physical assaults. To provide maximal physical protection, the robotic bodyguards need to consider the movement of the crowd as well as the obstacles in the environment. We propose two algorithms: *Threat Vector Resolution* (TVR) for a single bodyguard robot and *Quadrant Load Balancing* (QLB) for multiple bodyguards acting as a team. We evaluate the proposed approaches using metrics of the threat level and movement dynamics. The simulation study compares the results of the proposed approaches against rigid bodyguard team formations for various crowd configurations and team sizes.

Introduction

The capabilities of human-scale mobile robots are gradually reaching a level where robots can act in environments in the presence of dense populations of humans. Examples of practical applications include robotic museum guides (Thrun et al. 1999), people tracking (Schulz et al. 2003), telepresence (Michaud et al. 2007) and so on.

In this paper we are considering the problem of one or more robotic bodyguards providing effective physical cover to secure a human VIP from physical assaults. The VIP is moving in a public space in the presence of other humans. The robotic bodyguards must assess the members of the crowd as possible threats and position themselves accordingly. When there are multiple robotic bodyguards, they must act as a team, collaborating to minimize the level of threat. The robotic bodyguards must take into account the movement of the VIP, the density and movement pattern of the crowd as well as the natural obstacles and covers in the environment. We focus only on the direct physical assaults that can be protected against by providing physical cover (protection strategies against armed assailants or snipers are based on very different principles).

The expertise and training techniques of human bodyguards (“close protection operatives”) are clearly relevant to this problem. Unfortunately, the training manuals use narrative descriptions and in-situ practical examples, not directly applicable to the control of mobile robots.

In this paper, after introducing several metrics to quantify the threat to the VIP we propose two algorithms designed to lower the threat level. The *Threat Vector Resolution* (TVR) algorithm controls a single robotic bodyguard by performing probabilistic threat assessments and the resolution of multiple threat vectors. To extend this model to teams of bodyguard robots, we introduce the *Quadrant Load Balancing* (QLB) algorithm in where a team of robotic bodyguards share the load of physical cover task. In an experimental study, we compare the performance of these algorithms against rigid formation strategies.

Related Work

The robotic bodyguard problem is related to a number of research areas that received significant attention in recent years, such as robot team coordination, patrol scheduling, security checkpoint allocation and human crowd modeling.

Many studies related to security such as ARMOR (Pita et al. 2008), IRIS (Tsai et al. 2009), GUARDS (Pita et al. 2011) and RaPtoR (Varakantham, Lau, and Yuan 2013) consider the placement of checkpoints and deployment of patrol teams to provide protection against probable attacks by imminent adversaries. This requires generating mixed strategies for a group of defenders and adversaries using an exponential number of routes or schedules, which increases the computation requirements for the autonomous robots with limited computation, communication and power resources. (Khan, Arif, and Bölöni 2014) propose a technique in which robots learn to imitate human strategies to resolve the micro-conflicts that occur while moving in a dense crowd.

A multi-robot patrolling framework for cross-cultural environment is proposed in (Khan et al. 2012). The framework captures and analyzes the behavioral perception of the actions of the soldier and the robot by the local population. The multi-robot patrolling problem is usually formulated (Agmon, Kaminka, and Kraus 2011; Vanek et al. 2010) as a variation of the classical traveling salesman problem. Multi-agent based patrolling requires exponential decision making in order to minimize time lags between two consecutive visits of the agents to the same locations or gain advantage over adversary by protecting a particular geographical area.

Bodyguard Positioning Problem

We are considering a very important person (VIP) protected by a group of bodyguard robots $R = \{r_1, r_2, \dots, r_p\}$. The crowd members $G = \{g_1, g_2, \dots, g_k\}$, perform purposeful movement according to their own goals, but they also represent a threat of physical assault to the VIP. The robotic bodyguards aim to minimize the risk of physical harm to the VIP by providing physical cover: they position themselves to prevent potential attackers from reaching the VIP. To quantify the protection offered by a team of bodyguards, we first need to provide definitions to the concepts of *line-of-sight* and *safe distance*.

Definition 1 The line of sight function $LoS(x, y) \in \{0, 1\}$ specifies whether the agent x can observe the other agent y .

In the 2D space we are considering, $LoS(x, y) = 1$ implies that there is no obstacle between the positions of the agents. In defining line of sight, we assume that the robotic bodyguards have a 360° vision, thus their line of sight does not depend on the direction they are facing. This is enabled by the fact that there are many relatively inexpensive 360° degree cameras and laser scanners such as Hokuyo URG-04LX-UG01 (Tsui et al. 2012) currently being used on mobile robots.

Definition 2 The safe distance, *SafeDist*, is defined such that if $Dist(VIP, g_i) \geq SafeDist$ then crowd member g_i can be ignored in the assessment of threats to the VIP.

The safe distance captures the intuition that the robotic bodyguard would be able to physically stop an attacker before reaching the VIP if it starts outside the specified distance. The safe distance depends on the speed of the robot compared to a human, but in general it can be 10-15 meters or more. In crowded public places it is not feasible to keep every crowd member at a safe distance.

Definition 3 We define the threat level $TL(g_i, VIP) \in [0, 1]$ as the probability that crowd member $g_i \in G$ can successfully assault the VIP.

Human close protection operatives perform the assessment of the threat level based on physical, physiological and environmental attributes, taking advantage of their training, real life experience and intuition. Robotic bodyguards do not have the advantage of intuition and psychological assessment. In our model, we estimate the TL using the distance from the VIP as follows.

$$TL(g_i, VIP) = \begin{cases} LoS(g_i, VIP) \cdot e^{-(A \cdot (Dist(g_i, VIP))) / B} & \text{if } Dist(g_i, VIP) < SafeDist \\ 0 & \text{if } Dist(g_i, VIP) \geq SafeDist \end{cases} \quad (1)$$

where $Dist(g_i, VIP)$ is the distance between the crowd member and the VIP. A and B are the positive constants which define the slope and the magnitude of the risk curve respectively.

Definition 4 We define the cumulative threat CT posed to a VIP by a crowd $G = \{g_1, g_2, \dots, g_k\}$ as:

$$CT(G, VIP) = 1 - \prod_{i=1}^k (1 - TL(g_i, VIP)) \quad (2)$$

Note that CT is defined as a cumulative probability of the TL values. As the VIP and the crowd are moving, the cumulative threat level will change $CT = CT(t)$. Next, we need to find a metric that measures the threat to which a VIP is exposed during the full scenario.

Definition 5 Let us consider a VIP moving through a crowd over a scenario from time $t = 0$ to T . We define the total threat TT as:

$$TT = \int_0^T CT(t) dt \quad (3)$$

For a given crowd, the TT value depends on the speed of movement and the trajectory of the VIP. If the VIP moves faster and avoids dense crowds, the total threat will be smaller.

Let us now consider the ways bodyguards lower the threat level. We will call the *residual threat* and denote with $RT(g_i, VIP, R)$ the threat to the VIP assuming the presence of the bodyguard team $R = \{r_1, r_2, \dots, r_p\}$ at specific locations. This threat level depends on the relative location of the VIP, the threat and the bodyguards. In general, the threat level is lowered when the threat's possible ways for approaching to the VIP are blocked by a bodyguard.

Having defined the residual threat level RT , we define the cumulative residual threat $CRT(G, VIP, R) = 1 - \prod_{i=1}^k (1 - RT(g_i, VIP, R))$ and the total residual threat $TRT(R, C) = \int_0^T CRT(t) dt$ in analogy to the definitions of cumulative threat and total threat.

In addition to the goal of protection, the bodyguard robots must also operate discreetly and "smoothly". Frequent changes in the relative position of bodyguards distracts the VIP and increase the energy consumption. Taking a physical analogy, we aim to minimize the mechanical work performed by the bodyguard robots (that is, the curvilinear integral of the force $\vec{F} = m\vec{a}$ over their trajectory). The physical analogy is not perfect, because we do not really care about the mass of the bodyguard robots. Thus, the *total cumulative acceleration* TCA will be the integral over the absolute value of the accelerations summed for all the bodyguard robots:

$$TCA(R) = \sum_{i=1}^p \int_0^T |\vec{a}_i(t)| dt = \sum_{i=1}^p \int_0^T \left| \frac{d\vec{v}_i(t)}{dt} \right| dt \quad (4)$$

Frequent changes in the positioning decisions of the robotic agents produce higher TCA while it may or may not produce lower TRT. By the above definitions, the bodyguard positioning problem is having lower TRTs with acceptable TCAs for various scenarios.

Bodyguard Positioning Algorithms

Threat Vector Resolution

The threat vector resolution (TVR) algorithm is designed to position a single bodyguard robot such that the total threat is minimized. We define the threat vector $T\vec{V}$ as the sum of the unit vectors from the crowd members to the VIP, weighted

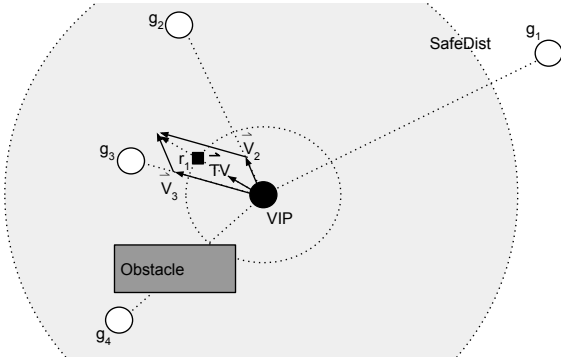


Figure 1: Positioning of a single robot r_1 with TVR.

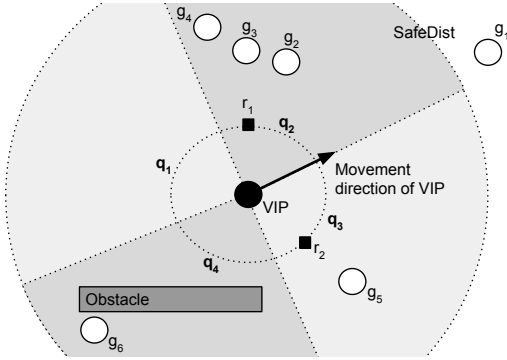


Figure 2: Assigning two robots to the quadrants q_2 and q_3 with highest loads.

by the threat level:

$$TV(\vec{VIP}, G) = \frac{\sum_{i=1}^k TL(g_i, VIP) \cdot \frac{\vec{V}(g_i, VIP)}{\|\vec{V}(g_i, VIP)\|}}{\sum_{i=1}^k TL(g_i, VIP)} \quad (5)$$

The threat level TL already encompasses the calculation of lines of sight due to obstacles as well as the maximum threat distance. According to the TVR algorithm the bodyguard robot will position itself at a fixed distance from the VIP in the direction of the threat vector.

Fig. 1 shows an example of the operation of the TVR algorithm. Crowd members g_1 and g_4 do not pose a threat to the VIP: g_1 is farther than *SafeDist* and g_4 is blocked by an obstacle. g_2 and g_3 create their threat vectors \vec{V}_2 and \vec{V}_3 respectively. As g_3 is closer to the VIP, its threat vector will be larger. The unit vector $T\vec{V}$ is obtained by summing and normalizing the threat vectors. The bodyguard robot r_1 positions itself at a constant distance from the VIP in the direction of $T\vec{V}$.

Quadrant Load Balancing

TVR is not appropriate for the control of teams of bodyguards because it will position all the robots very close to each other in the direction of the threat vector, ignoring threats coming from other directions. To overcome this, we propose the *Quadrant Load Balancing* (QLB) algorithm that

Algorithm 1 Quadrant Load Balancing (QLB)

```

procedure QLB( $G, VIP, R$ )
   $R \leftarrow R - r_i$ 
  while  $Q \neq null$  do
    Select  $q_{max} \in Q$  with max  $l_q$ 
    if  $q_{max} = null$  then
       $q_{max} \leftarrow r_i \cdot q_{default}$ 
    end if
     $w_{max} \leftarrow \|g_{q_{max}}\|$ 
    for all  $r \in R$  do
      if  $r.q = q_{max}$  then
        if  $r.w < w_{max}$  then
           $w_{max} \leftarrow w_{max} - r.w$ 
        else
           $w_{max} \leftarrow 0$ 
        end if
      end if
    end for
    if  $w_{max} > 0$  then
       $\vec{p} \leftarrow TV(g_{q_{max}}, VIP)$ 
       $P_{r_i} \leftarrow PXY(\vec{p})$ 
      return
    else
       $Q \leftarrow Q - q_{max}$ 
    end if
  end while
end procedure

```

allows a team of bodyguards to more evenly distribute the tasks of protection from specific crowd members. In this approach, the protection circle of a VIP is divided into quadrants with each bodyguard being responsible for protecting against threats in one or more quadrants. This protection can be more or less difficult depending on the number of threats in the quadrant and their threat level TL. The intuition behind the QLB algorithm is that the bodyguards must have tasks with approximately equal difficulty (their load must be balanced).

Algorithm 1 describes the operation of the QLB model. $Q = \{q_1, q_2, q_3, q_4\}$ is the set of quadrants, G_{q_i} is the set of crowd members corresponding to the quadrant q_i . $R = \{r_1, r_2, \dots, r_k\}$ represents the set of robotic bodyguards except r_i who generate QLB request, where k is initially equal to the total number of robot bodyguards. $L = \{l_{q1}, l_{q2}, l_{q3}, l_{q4}\}$ is a set of threat level associated with individual quadrant (q). The algorithm iteratively evaluates all the quadrants in the decreasing order of quadrant load. The workload associated with robot bodyguard r is given by $r.w$ and w_{max} represents the load value of quadrant with current maximum load q_{max} .

$PXY(\vec{p})$ is the proximity function that provides an unoccupied position close to the position vector \vec{p} . The position vector \vec{p} is the resolved vector over the sum of the threats in the quadrant q computed by Equation 5. PXY returns the

position as follows:

$$PXY(\vec{p}) = \begin{cases} P_{VIP} + \vec{p} \cdot ProtectDist & \text{if unoccupied} \\ PXY((\vec{p})') & \text{otherwise} \end{cases} \quad (6)$$

where

$$(\vec{p})' = (\vec{p}.x \cdot \cos(\theta), \vec{p}.y \cdot \sin(\theta)), \quad (7)$$

such that $ProtectDist$ is the protection distance of the robot bodyguards from the VIP. $ProtectDist$ is shown by the inner circles in Fig. 1 and Fig. 2. θ is the minimum angle which produces an unoccupied location value in any one of the two directions for $PXY((\vec{p})')$. We define a location as occupied if there exists another robot bodyguard or an obstacle on it. Hence, PXY function produces an unoccupied location having the same distance from P_{VIP} and being closest to the ideal case of $P_{VIP} + \vec{p} \cdot ProtectDist$.

Algorithm 1 assigns the quadrant with the highest TL to q_{max} at each iteration of the while loop. In case of no threat in any quadrant, the bodyguard r_i who is running QLB follows a locker-room agreement (pre-decided before execution) by making q_{max} as $r_i \cdot q_{default} \cdot q_{default}$ is the default quadrant which is assigned to the bodyguard r_i before starting the operation. In the inner loop, the algorithm checks for the total load of the quadrant q_{max} covered by existing bodyguards in the same quadrant. If the load of the quadrant q_{max} is higher than the total load covered by existing robot bodyguards, then the bodyguard r_i move to the position on q_{max} given by P_r , otherwise, this process repeats for the rest of the quadrants.

Fig. 2 illustrates an example case of the QLB algorithm. In this figure, the quadrant q_1 has no crowd member, q_4 has no visible crowd member ($l_{q_3} = l_{q_4} = 0$) and $q_{max} = q_2$. Among the two unassigned robotic bodyguards, r_1 is first assigned to q_2 with the highest load. In the second iteration, the load of q_2 is updated and r_2 is assigned to the quadrant q_3 with the highest remaining load l_{q_3} .

Simulation Study

Simulation environment

We carried out simulation experiments using the Java-based YAES Simulator (Bölöni and Turgut 2005) developed by our research group. The simulator is interconnected with the V-Rep simulator from Coppelia Robotics. Our simulations do not use the physics engines from V-Rep, but they use simple control models.

We compared the experimental results of four crowd configurations: “static crowd”, “with the flow”, “against the flow” and “mixed crowd”. In the “static crowd” configuration, crowd members are static and they occupy distinct locations on the map. This configuration provides the base case for comparing sudden changes in the threat metric with respect to crowd movement. In the “with the flow” configuration, crowd members move in the same direction as the VIP, often encountered in situations where the crowd and the VIP aim to enter or exit a building. In this configuration the robotic bodyguards need to evaluate the threats approaching from the rear. In the “against the flow” configuration,

the crowd moves against the direction of the VIP. Similarly, in the “mixed crowd” configuration, each crowd member chooses a random destinations on the map and traverses the shortest path toward the selected destinations using the D* Lite algorithm (Koenig and Likhachev 2002). In all configurations, the crowd members initialize with random movement speeds between 1ft to 5ft per time step, before the start time of simulation. These variations in speed represent diversity of gait preferences of the humans in crowd. In the simulation, the VIP traverses through 10 different paths on a given map and configuration covering most possible cases of interactions of crowd and robotic bodyguards. Each simulation run corresponds to a different path taken by the VIP on the simulation map.

Performance Results

Swarm robotics algorithms focus on maintaining fixed team formations of robots while performing path planning. We compared our approaches against the fixed formation in order to assess the decrease in TRT of the configuration. Moreover, we compared effectiveness of the approach over different number of robotic bodyguard agents deployed in the simulation. Realistic bodyguard training literature provides these fixed formations which depend on various factors such as location, threat intensity, number of bodyguards, and so on. For the fixed formations, we considered Cartesian plane-based quadrants where origin and ordinate axis of the plane as VIP coordinates and heading direction, respectively. The order of placement varies as the number of bodyguards increases.

Experiment 1: Cumulative Residual Threat We first compared the cumulative threat CT and cumulative residual threat CRT values at every accessible location for the VIP over the simulation map with static crowd configuration as shown in Fig. 3. Fig.3-middle shows distribution of CT values in the case of no protection. The peaks represent the locations with lowest security, which correspond to the locations of the crowd members in Fig.3-left. Flat regions in Fig.3-middle correspond to the obstacles and walls. Fig. 3-right reveals the effect of the QLB algorithm with 3 bodyguards over the same crowd configuration. We observed that QLB produces significantly lower CRT values in most regions of the map compared to the case of no protection. However, peaks still exist due to the fact that if the VIP is too close to a crowd member, blocking an attack becomes almost impossible by a robotic bodyguard. On the other hand, the goal of the proposed positioning algorithms is avoiding these cases by preventing the crowd members to stay in the protection circle of the VIP.

Experiment 2: Total Residual Threat In this experiment, we compared the total residual threat TRT values of seven bodyguard strategies over multiple simulation runs in four crowd configurations. The TRT for the ‘No bodyguard’ case is actually the same as the total threat TT value, and illustrates the actual threat to the VIP from the crowd members in the absence of bodyguards. We are considering these values as benchmark for comparison of the TRT values. We did not include the results for QLB with 1 bodyguard since




Figure 3: Comparison of CT and CRT over simulation map. Left: static crowd configuration, Middle: CT for no bodyguard, Right: CRT for 3 bodyguard robots with QLB.

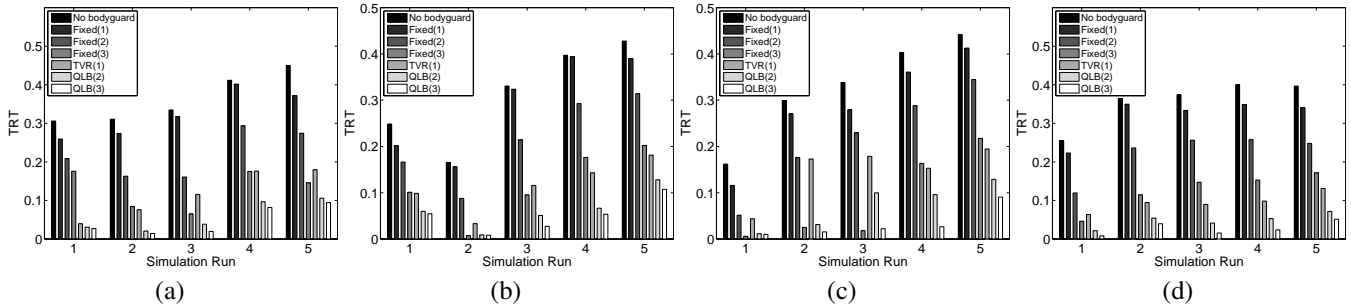


Figure 4: TRT comparison of No bodyguard, 1 bodyguard, 2 bodyguards and 3 bodyguards with fixed formation, TVR and QLB. The different crowd configurations were (a) static crowd, (b) “with the flow”, (c) “against the flow” and (d) mixed.

QLB produces very similar output to TVR in that case. In general, the lower the TRT, the better the VIP is protected in the given scenario.

We start with the “static crowd” configuration in Fig. 4-(a). In this scenario, the differences in the TRT values of the ‘No bodyguard’ scenario over multiple simulation runs depend on the movement of the VIP - for instance, sometimes the presence of an obstacle lowers the TRT by providing physical cover from crowd. As expected, the higher the number of robotic bodyguards, the lower the TRT values. Nevertheless, the proposed adaptive strategies provide a better protection than the fixed formation strategies. For instance, the TVR strategy with a single bodyguard provides protection comparable to and sometimes (simulation runs 1 and 2) significantly better than the fixed formation strategy with 3 bodyguards. As expected, the QLB strategy with 3 bodyguards provided the lowest TRT for all simulation runs.

Fig. 4-(b) shows the results of the “with the flow” configuration, where the crowd is moving in the same direction as the VIP, although their speeds and final destinations may differ. In the fixed configuration, the bodyguards position themselves at behind-left and behind-right of the VIP at an arm length distance. Therefore, the VIP has less frequent variation in the TRT in case of approaching crowd from behind. Due to this reason, 3 robotic bodyguards sometimes perform better as shown in the second simulation run.

Fig. 4-(c) shows the results of the “against the flow” configuration, where the crowd is moving in the opposite di-

rection from the VIP. The increase in the TRT values of ‘No bodyguard’ is due to the increase in the crowd density and the movement time. 3 bodyguards, with fixed formation strategy ‘Fixed(3)’, also slow down to protect VIP as the crowd density increases. 3 bodyguards with QLB algorithm ‘QLB(3)’ consistently produces lower TRT values for all the simulation runs.

Finally, Fig. 4-(d) reveals the results of the “mixed crowd” configuration. As in the other scenarios, QLB with 3 bodyguards is the winner as it produces lower TRT values while TVR with 1 bodyguard performs better than fixed formation with 3 bodyguards for all distinct VIP paths. Overall, we observed that TVR and QLB provide lower TRT values for various movement paths of the VIP and the crowd configurations.

Experiment 3: Mean Total Cumulative Acceleration

In this experiment, we presented Mean-TCA results of different bodyguard teams deploying fixed and algorithmic strategies for multiple simulation runs as shown in Fig. 5. Well-coordinated bodyguard teams should have minimal Mean-TCA with the low Mean-TRT values at all times. The higher Mean-TCA values are caused by various factors such as frequent shuffling of positions, one agent performing more work than the other team members or noisy communication in threat assessment and task allocation.

As shown in Fig. 5, fixed formation strategies have better coordinated movement compared to the algorithmic strategies deployed by the bodyguard teams because of the un-




Figure 5: Mean-TCA comparison of 1 bodyguard, 2 bodyguards and 3 bodyguards with fixed formations, TVR and QLB for different crowd configurations.

affected behavior of fixed formation team by crowd movement. In static crowd configuration, Mean-TCA attained by ‘TVR(1)’ is very close to Mean-TCA by 3 bodyguards team in fixed formation. This is a result of no abrupt changes in threat to VIP in the absence of crowd movement. Moreover, the Mean-TCA results of 2 bodyguards team in “with the flow” and “against the flow” configurations are much higher compared to the results of other bodyguard teams, due to the frequent shuffling in alternate quadrants by 2 bodyguards during parallel and opposite crowd movements. In “with the flow” configuration, 3 bodyguards team has equal Mean-TCA to 1 bodyguard team. This is due to one bodyguard performing excessive task of guarding multiple quadrants compared to other bodyguards in the team. The excessive increase in the Mean-TCA of multiple bodyguards in the “mixed crowd” configuration is a result of frequent miscommunication and misjudgment in identifying each others’ position and intentions. Fig. 5 reveals the improvement required on QLB algorithm in terms of the effective communication among bodyguards.

Conclusions

In this paper, we focused on the positioning of single and multiple robotic bodyguards during movement to protect the VIP. We proposed *Threat Vector Resolution* (TVR) approach for single robot bodyguard positioning and *Quadrant Load Balancing* (QLB) for collaborative security using multiple robot bodyguards. We evaluated the proposed approaches against rigid formation of robotic bodyguards by the simulation experiments. The proposed approaches provide better total residual threat values for various movement scenarios and bodyguard teams with different sizes.

Acknowledgements: The research reported in this document/presentation was performed in connection with Contract Number W911NF-10-2-0016 with the U.S. Army Research Laboratory. The views and conclusions contained in this document/presentation are those of the authors and should not be interpreted as presenting the official policies

or position, either expressed or implied, of the U.S. Army Research Laboratory, or the U.S. Government unless so designated by other authorized documents. Citation of manufacturer’s or trade names does not constitute an official endorsement or approval of the use thereof. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

- Agmon, N.; Kaminka, G. A.; and Kraus, S. 2011. Multi-robot adversarial patrolling: Facing a full-knowledge opponent. *Journal of Artificial Intelligence Research* 42:887–916.
- Bölöni, L., and Turgut, D. 2005. YAES - a modular simulator for mobile networks. In *Proc. of the ACM/IEEE MSWIM’05*, 169–173.
- Khan, S. A.; Arif, S.; and Bölöni, L. 2014. Towards learning movement in dense crowds for a socially-aware mobile robot. In *Workshop on Adaptive Learning Agents (ALA-2014)*.
- Khan, S. A.; Bhatia, T. S.; Parker, S.; and Bölöni, L. 2012. Modeling human-robot interaction for a market patrol task. In *Proc. of the FLAIRS’12*, 50–55.
- Koenig, S., and Likhachev, M. 2002. Improved fast replanning for robot navigation in unknown terrain. In *Proc. of the IEEE ICRA’02*, volume 1, 968–975.
- Michaud, F.; Boissy, P.; Labonte, D.; Corriveau, H.; Grant, A.; Lauria, M.; Cloutier, R.; Roux, M.-A.; Iannuzzi, D.; and Royer, M.-P. 2007. Telepresence robot for home care assistance. In *AAAI Spring Symposium: Multidisciplinary Collaboration for Socially Assistive Robotics*, 50–55.
- Pita, J.; Jain, M.; Marecki, J.; Ordez, F.; Portway, C.; Tambe, M.; Western, C.; Paruchuri, P.; and Kraus, S. 2008. Deployed ARMOR protection: the application of a game theoretic model for security at the Los Angeles International Airport. In *Proc. of the AAMAS’08*, 125–132.
- Pita, J.; Tambe, M.; Kiekintveld, C.; Cullen, S.; and Steigerwald, E. 2011. GUARDS: Innovative application of game theory for national airport security. In *Proc. of the IJCAI’10*, volume 22, 2710.
- Schulz, D.; Burgard, W.; Fox, D.; and Cremers, A. B. 2003. People tracking with mobile robots using sample-based joint probabilistic data association filters. *Int’l Journal of Robotics Research* 22:99–116.
- Thrun, S.; Bennewitz, M.; Burgard, W.; Cremers, A. B.; Dellaert, F.; Fox, D.; Hahnel, D.; Rosenberg, C.; Roy, N.; Schulte, J.; et al. 1999. MINERVA: A second-generation museum tour-guide robot. In *Proc. of IEEE ICRA’99*, volume 3.
- Tsai, J.; Kiekintveld, C.; Ordonez, F.; Tambe, M.; and Rathi, S. 2009. IRIS - a tool for strategic security allocation in transportation networks. In *Proc. of AAMAS’09*, 37–44.
- Tsui, K. M.; Von Rump, S.; Ishiguro, H.; Takayama, L.; and Vicars, P. N. 2012. Robots in the loop: telepresence robots in everyday life. In *Proc. of the ACM/IEEE HRI’12*, 317–318.

Vanek, O.; Bosansky, B.; Jakob, M.; and Pechoucek, M. 2010. Transiting areas patrolled by a mobile adversary. In *Proc. of the IEEE CIG'10*, 9–16.

Varakantham, P.; Lau, H. C.; and Yuan, Z. 2013. Scalable randomized patrolling for securing rapid transit networks. In *Proc. of the IAAI'13*, 1563–1568.