Agents vs. Pirates: Multi-Agent Simulation and Optimization to Fight Maritime Piracy

Michal Jakob, Ondřej Vaněk, Ondřej Hrstka and Michal Pěchouček
Agent Technology Center, Faculty of Electrical Engineering, Czech Technical University in Prague
Technicka 2, Praha 6, Czech Republic
{jakob, vanek, hrstka, pechoucek}@agents.fel.cvut.cz

ABSTRACT

Contemporary maritime piracy presents a significant threat to the global shipping industry, with annual costs estimated at up to US$12bn. To address the threat, commanders and policymakers need new data-driven decision-support tools that will allow them to plan and execute counter-piracy activities most effectively. So far, however, the provision of such tools has been very limited. To fill this gap, we have employed the multi-agent approach and developed a novel suite of computational tools and techniques for operational management of counter-piracy operations. A comprehensive agent-based simulation enables the stakeholders to assess the efficiency of a range of piracy counter-measures, including recommended transit corridors, escorted convoys, group transit schemes, route randomization and navy patrol deployments. Decision-theoretic and game-theoretic optimization techniques further assist in discovering counter-measure configurations that yield the best trade-off between transportation security and cost. We demonstrate our approach on two case studies based on the problems and solutions currently explored by the maritime security community. Our work is the first integrated application of agent-based techniques to high-seas maritime security and opens a wide range of directions for follow-up research and development.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multi-agent Systems; I.6.3 [**Simulation and Modeling**]: Applications

General Terms

Algorithms, Security, Experimentation

Keywords

agent-based modeling, simulation, transportation, maritime piracy, security, optimization, case study, operation research

1. INTRODUCTION

Contemporary maritime piracy presents a significant threat to the global shipping industry, with annual total costs estimated at up to US$12bn [3]. International efforts have led to the reduction of the success rate of pirate attacks. The number of pirate attacks, average amount of ransom paid and the number of crewmembers held in captivity, however, remain high. In the first nine months of 2011 and for Somali-based piracy alone, there were 194 attacks reported for Somalia, 24 vessels hijacked and 400 crewmembers taken hostage[1].

From the many levels on which solutions to the problem are sought, we focus on the operational management of the situation at sea, as this is the arena where progress can be made in the short term, before long-term sustainable solutions can be developed onshore. Operational management is also the area where the agent-based approach can provide significant added value—due to the lack of strong central authorities, complex structure of shared, competing and hostile interests and fragmented, distributed nature, global maritime shipping is best viewed as a complex multi-agent system and, consequently, solutions pursued within the multi-agent framework.

In line with the efforts of main governmental and industry stakeholders, we focus on planning and management of operational piracy counter-measures designed to increase the security of maritime transit, including recommended transit corridors, group transit and escorted convoy schemes, route randomization and coordinated patrol deployments. When properly designed and implemented, these measures have a potential to significantly improve transportation security with reasonable additional cost. However, due to complex spatial and temporal dependencies between individual counter-measures, discovering their effective, synergistic combinations presents a major challenge.

We address this challenge in a novel way by a combination of agent-based modeling and optimization techniques. We have built a data-driven piracy-aware agent-based model of maritime transportation, which enables decision makers to reduce uncertainty about the effects of their management and regulatory actions. The model integrates a wide range of real-world data and, to our best knowledge, is the first computational model that simulates global shipping down to a level of individual vessels. This is crucial for accurately capturing emergent, collective effects that arise from the coordination and cooperation of commercial and navy vessels and their non-cooperative interaction with pirates.

Recognizing that being able to assess effects of piracy counter-measures is only a first step towards discovering their most effective combinations, we have also developed computational optimization techniques that partially auto-

[1]Source: IMB Piracy Reporting Centre

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems – Innovative Applications Track (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

mate counter-measure design. We demonstrate the joint utility of the simulation and optimization tools on two case studies based on the problems and solutions currently explored by the maritime security community.

Please note that due to the wide scope of our research and the number of techniques developed, we only provide here a high-level description of key concepts, methods and components, with the prime focus on elucidating how the multi-agent perspective has been beneficial in addressing the problem at hand—references to papers studying individual issues in depth are provided where appropriate.

2. MARITIME TRANSPORTATION MULTI-AGENT SYSTEM

Maritime piracy takes place within a larger *maritime transportation system* comprised of all seaborne vessels engaged in maritime activities. At the level of abstraction suitable for operational management, the maritime transportation system can be viewed as a multi-agent system, with vessels corresponding to autonomous agents. Vessels in the system are capable of moving freely within the spatial boundaries of surface waters while interacting with the maritime environment, other vessels (either directly via communication or indirectly via environment) and other non-vessel actors (such as shipping operators or traffic coordinators). For most of the time, each vessel pursues its individual goals but there are also situations where multiple vessels interact. Such interactions are either non-cooperative (such as pirate attacks or navy warship pirate interceptions) or cooperative (such as merchant vessels' calls for help to navy warships).

2.1 Vessel Agents

The following types of vessels play the most important role in the dynamics of maritime piracy:

- *merchant ships* – large ocean-going vessels carrying cargo over long distances between world's ports; primary targets of pirate attacks;

- *pirate ships* – vessels of different types and sizes operating within and in close proximity to main shipping lanes, where they attempt to attack, board and hijack passing merchant vessels; depending on their operational area, pirate ships range from small skiffs up to large motherships acting as floating bases from which speedboats are launched to attack[2];

- *navy warships* – military vessels of different categories operating in piracy-affected areas and capable of armed action against pirates.

2.2 Piracy Counter-Measures

Merchant and navy vessels can participate in a range of piracy counter-measures designed to increase the security of voyage through piracy-affected waters. Most measures require cooperation between multiple vessels and can be viewed as multi-agent coordination mechanisms that complement standard, single-agent vessel behaviors. Based on discussions with the maritime security community, we consider the following operational piracy counter-measures:

[2]The group of one mothership and multiple accompanying speedboats is referred to as a *pirate attack group(PAG)*

Counter-measure	Parameters
Transit corridor	sequence of GPS waypoints
Patrol deployment	patrol stationary location and/or dynamic patrolling policy
Group transit	corridor, speed levels, transit schedule (per speed level)
Escorted convoy	corridor, departure time, convoy speed
Route randomization	corridor, randomization distribution

Table 1: Piracy counter-measures considered and sets of parameters by which they are specified.

- *Recommended transit corridors* concentrate merchant traffic along defined routes connecting sequences of waypoints. Corridors facilitate protection from navy vessels; however, they also makes targeting transiting vessels easier for pirates.

- *Group transit schemes* coordinate the timing of merchant vessel transit so that vessels pass high-risk piracy areas in groups; this improves mutual awareness and facilitates navy response; however, it makes the transit take longer as vessels have to follow a predefined entry schedule and may have to reduce their cruising speeds to match the speed of their respective transit group.

- *Patrol deployments* position navy warships in strategic locations from where they can provide assistance to nearby merchant vessels in case of a pirate attack; we consider both stationary and dynamic deployments. Navy patrols are very effective locally; however, their action radius is limited and their sustained operation carries huge costs.

- *Escorted convoy schemes* arrange incoming merchant vessels into escorted convoys prior to transiting a high-risk area. In contrast to the group transit scheme, the convoy is accompanied by a dedicated armed escort vessel throughout the transit. Although very effective, a large-scale convoy system would require a very high number of escort vessels to operate effectively, which is costly.

- *Route randomization* relaxes transit corridor boundaries by making transiting vessels randomly deviate from the center of the corridor according to a given probability distribution (typically uniform). Route randomization reduces predictability of vessel positions, and thus makes planning and execution of pirate attacks more difficult.

Each counter-measure can be parameterized by a set of parameters (see Table 1). Except for route randomization, all above measures are currently actively used, although convoy schemes are operated rather sporadically by national navies on an ad-hoc basis. The use of transit corridors and group transits is currently limited to the Gulf of Aden.

2.3 System Parametrization

From the perspective of maritime security, there are several factors fundamentally affecting temporal and geographical distribution of pirate attacks:

- *Merchant traffic patterns* – represented in terms of an *origin-destination matrix* describing a yearly number of trips between world's major shipping ports.

- *Pirate population* – represented by the number of active pirate groups and locations of their on-shore bases.

- *Weather conditions* – represented by spatio-temporal maps of wind direction and speed, sea state and visibility.

- *Piracy counter-measures* – represented by a list of specific counter-measures and their parameters.

The above factors comprise a minimum set of parameters that has to be specified for a maritime transportation system before its piracy-related properties can be studied, both in the real-world and in a simulation.

2.4 Performance Metrics

A wide range of events and measurable quantities can be observed on a maritime transportation system. On the system level, the following metrics are of practical interest with regards to evaluating and optimizing the efficiency and security of maritime shipping:

- *pirate attack statistics* – the number of pirate attacks that occurs in the system in a defined time period; we further distinguish between *successful attacks* (=hijacks), *intercepted attacks* (attacks that fail due to navy interception) and *aborted attacks* (attacks aborted by pirates themselves, often due to effective employment of self-defense measures by the targeted vessel);

- *average transit distance* – average distance travelled per merchant vessel trip (in kilometers);

- *average transit duration* – average duration of merchant vessel trip (in hours).

Additional metrics can be derived from the primary metrics, such as fuel consumption or operational cost per trip.

2.5 Problem Statement

Given the framework introduced above, the problem addressed by our work can be more precisely defined as:

1. Analyzing relationships between the parameters of the global maritime transportation system and its performance metrics

2. Discovering such combinations of piracy counter-measures that optimize a user-defined function over the performance metrics (i.e. a user-defined trade-off between the security and cost of shipping)

In the following two sections, we show how we addressed both problems.

3. SIMULATION

In order to address the first problem, we have built an agent-based model/simulation[3] of the global maritime transportation system. The simulation closely follows the multi-agent conceptualization of the transportation system described in the previous section. As such, it provides *executable* models of vessel behaviors as well as the (collective) piracy counter-measure.

Given the critical role the interactions between merchant, pirate and navy vessels play in the dynamics of maritime piracy, agent-based, micro-simulation approach is vital for accurately modeling the effect of piracy on maritime transportation, as it allows to capture phenomena and provide detail of analysis not attainable with macro-level equation-based methods [15].

3.1 Input Data

To achieve a sufficient level of accuracy, data-driven agent-based modeling requires large amounts of data for setting, calibrating and validating individual parts of the model. In contrast to macro-modeling approaches, data both on individual and macro level are required. To build the model of global maritime transportation, we have used the following categories of data: (1) *geographical data* (shore lines, shallow waters), (2) *weather data* (visibility range, sea state), (3) *merchant traffic data*[4] (origin-destination matrix, fleet composition, vessel operational characteristics), (4) *pirate intelligence* (base locations, attack strategies, capabilities, historic reported pirate incidents[5]), (5) *military operations* (number of warships and their operational areas) and (6) *piracy counter-measures* (see the previous section).

Obtaining the above data in quantity and quality required was and remains to be a major challenge due to enormous fragmentation of data gathering activities in the maritime domain and commercial sensitivity of some of key data. Moreover, many data sets obtained were noisy and incomplete and required significant preprocessing before integration to the model.

3.2 Agent Behavior Implementation

Maritime transportation simulation requires agent control architecture capable of expressing desired individual and collective vessel behaviors. Agents have to be able to execute long-running actions while reacting to interruptions. The minimum intelligent agent architecture that can handle such requirements is a model-based reflex agentwith an encapsulated deliberative module handling route-planning. The required class of behaviors should be implementable in a modular and extensible way, facilitating sharing of common behavior fragments between different classes of vessels. At the same time, the agent control architecture should be computationally efficient enough to handle thousands of simulated agents. Unfortunately, none of existing agents architectures or simulation platforms supports these requirements. General agent-based simulation toolkits do not provide sufficient abstractions for modular behavior implementation; such abstractions are provided by cognitive agent architectures (e.g. Jazzyk [10]) but these require substantial computational resources to run.

Our implementation therefore uses *extended finite state machines (FSM)*, which augment standard FSMs with internal variables associated with each state. Individual states correspond to the principal mental states of the vessel agent

[3]We use both terms rather interchangeably, choosing *model* when focusing on the descriptive aspect and choosing *simulation* to emphasize the dynamic/execution aspect

[4]Merchant vessel trajectories in the form of AIS records are available at e.g. AISLive: `http://www.aislive.com/`

[5]Pirate reports are available from e.g. OceanusLive: `http://www.oceanuslive.org`

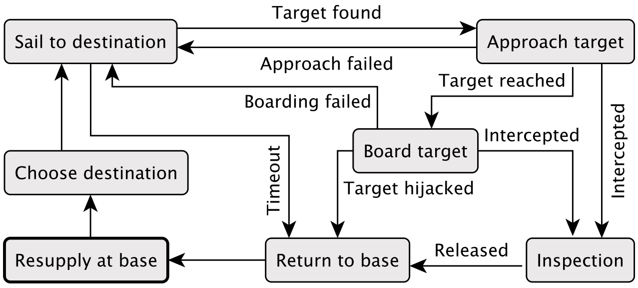

Figure 1: A finite state machine of the pirate vessel agent.

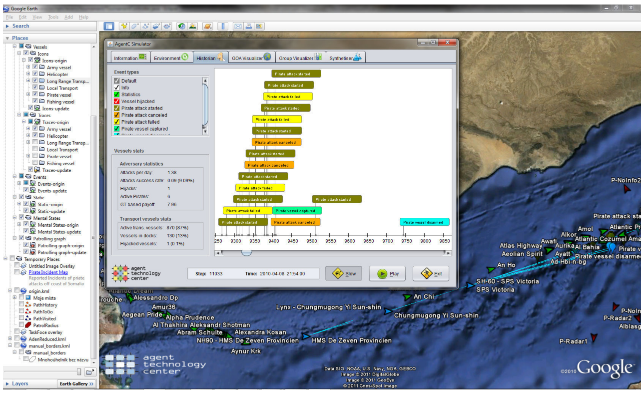

Figure 2: AgentC simulation platform—visualization of the Gulf of Aden group transit.

(such as move, attack, patrol etc.) and their associated actions. These actions may involve complex deliberative procedures, e.g. risk-aware route planning for merchant ships and adaptive target selection for pirates, giving the FSMs reasoning capabilities beyond simple reactive control. Although limited (e.g. not capable of executing concurrent activities), extended FSMs proved to provide a good trade-off between expressivity/modularity and computational efficiency. An example FSM is given in Figure 1.

3.3 Simulation Platform Implementation

Extended FSM-based behavioral models are executed on a Java-based multi-agent platform built partially using the lightweight Alite [6] multi-agent simulation toolkit and employing Google Earth for geo-spatial visualization (see Figure 2). The simulation platform provides abstractions for representing the maritime environment, agent-to-environment sensor interfaces and agent-to-agent communication protocols. Time-stepped simulation execution model is used, although we consider transitioning to the event-based model to further improve computational efficiency. Each simulation run is defined by a scenario defining the parameters of the maritime transportation system (Section 2.3). Parallel execution of large numbers of simulations is supported using the Eucalyptus cloud-computing platform[7].

3.4 Model Validation

The model has been validated both on the individual and system level. Individual-level models have been validated against independent test datasets capturing real-world behavioral patterns for the respective type of vessels. For example, merchant vessel models have been validated against real-world trajectories obtained from satellite AIS and voluntary reporting systems. In addition, system-level behavior has been validated against empirical data. The spatial distribution of merchant traffic has been compared with maritime shipping density maps. To validate the interaction of all types of vessels and counter-measures employed, we compared pirate incidents generated by the model with real-world piracy incidents. More details about the validation can be found in [17].

[6] `http://agents.fel.cvut.cz/projects#alite`
[7] `http://www.eucalyptus.com/`

4. OPTIMIZATION

The maritime transportation simulation introduced in the previous section empowers policy-makers in estimating the effects of different combinations of piracy counter-measures. Given the number parameters and combinations of these counter-measures, however, finding their right configuration remains difficult. We have therefore explored ways to provide computational support for optimizing counter-measure configurations automatically.

In its full generality, looking for such optimum configurations is a massive multi-agent optimization problem: a multi-objective performance function is optimized in a stochastic and partially observable environment with a high degree of uncertainty and thousands of interacting, largely self-interested agents employing a wide range of parameterized strategies and policies in the presence of multiple adaptive adversaries. Even if solutions in a form of a massive joint transit and patrolling routes and schedules could be found, they might be too complex and unstructured to be understood and hence trusted by human stakeholders. Instead of trying to solve the full problem, we therefore focused on optimizing individual counter-measures. For computational reasons, proposed optimization algorithms work with highly abstracted problem representations not containing the same amount of details as the simulation model. The simulation is therefore used to validate optimization results and to obtain higher-accuracy assessment of their expected real-world performance, which can be subsequently used to fine-tune proposed solutions.

4.1 Group Transit Optimization

Since participation in group transit schemes is voluntary, the problem of determining optimum speed levels and transit schedules can be viewed as a cooperative game with non-transferable utilities. Because of computational intractability of solving such a game for real-world problem sizes, we have so far considered two simplified, cooperative formulations of the problem.

The simpler formulation concerns the optimization of *fixed-schedule* group transit schemes. Taking into account the distribution of cruising speeds of transit traffic, we look for such a fixed set of speed levels that result in a shortest average transit duration. The optimum set of speed levels is found by searching for an optimum binning of the histogram of transit cruising speeds using a branch&bound search combined with dynamic programming (see [6] for details).

The more advanced formulation explores the potential of *dynamic* group transit schemes, in which speeds and schedules are not fixed in advance but determined on-the-fly, using multi-agent coordination techniques, based on incoming traffic. The added flexibility allows the dynamic group transit scheme to achieve higher performance than fixed-schedule schemes on the expense of more extensive vessel coordination and information sharing.

4.2 Randomized Transit Routing

A major disadvantages of fixed transit corridors is the high predictability of vessel positions [13], which makes targeting merchant traffic easier for pirates. Predictability can be reduced by instilling a certain amount of randomness in transit routing. A basic approach applies uniformly or normally distributed randomization to disperse the traffic away from the corridor center and/or to alternate between several predefined corridors. Better route randomizations can be obtained using game-theoretic techniques, which explicitly account for the payoff the merchant vessels and attacking pirates receive from different transit routes. To this end, we extended the model of security games and formalized the hostile area transit problem as a zero-sum normal-form game between two mobile players, the transit and the pirate, each choosing a route maximizing its utility. The solution, found using incremental equilibrium search techniques, can achieve up to two-fold reduction in the attack rate compared to the basic approach. Details are provided in [16, 18].

4.3 Optimum Patrol Deployment

Due to their limited numbers and very high operational costs, military vessels have to be deployed in a way maximizing their protective effect. We considered two formulations of patrol deployment problem. The basic formulation considers a *stationary* deployment of patrol vessels—each vessel is assigned a fixed location from which it only departs to assist vessels in danger. Given a number of patrol vessels, we look for such a set of deployment locations that maximizes the volume of commercial traffic within the patrol's effective action radius[8]). Deployment locations are found using a vector quantization *Lindo-Buzo-Gray (LBG) algorithm* [9]. See Figure 3 for an illustrative example.

The basic formulation does not take into account the ability of pirates to adapt their attack locations in response to their observation of fixed patrol deployments. To overcome this problem, we therefore considered a game-theoretic formulation of the problem in which patrols are mobile and randomize their movement to minimize the ability of pirates to take advantage of their predictable absence, while taking into account transit routes of individual merchant vessels. An novel extended-form game in the Stackelberg setting is used as the underlying formal model. The approach is particularly effective in combination with randomized transit routing. See [2] for more details.

5. CASE STUDIES

We have applied the developed modeling and optimization tools to several real-world use cases, based largely on feedback from the maritime community and discussions with officials from the International Maritime Organization and U.S. Office of Naval Research. Here we present two specific case studies—one focusing on evaluating the combination of a novel corridor system and patrol deployments in the piracy-infested Indian Ocean, the other on the optimization of group transit scheme in the Gulf of Aden. More information about the developed tools and their applications can be found at the AgentC project website[9].

[8]Radius of 150km is currently considered; this corresponds to the ability to respond within 40 minutes using an onboard helicopter.

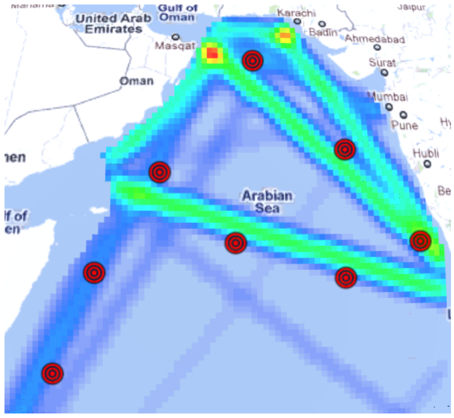

Figure 3: Example traffic density and the corresponding traffic coverage-maximizing deployment of 8 patrols.

Except for the parameters explicitly mentioned as variables of the study, the maritime system configuration remains the same throughout both case studies, in particular the origin-destination matrix capturing global merchant shipping flows. The simulation contained approximately 4500 merchant vessel agents, up to 100 navy warships and up to 20 pirate ship agents. In both case studies, the results given are for one year of simulated maritime traffic. Because parts of agent decision making are inherently nondeterministic, each configuration was simulated for 100 runs and average values are presented. One simulation run took approximately 10 mins of single 2.5GHz CPU core execution time.

5.1 CS I: Indian Ocean Corridor System

The *International Recommended Transit Corridor (IRTC)*, established in 2009, has since proven—in combination with the deployment of navy patrols—a very effective tool for suppressing successful pirates attacks in the Gulf of Aden. Recently, the maritime security community has been discussing the possibility of establishing additional corridors in the Indian Ocean, where pirate activity is also high following pirates' displacement from the Gulf of Aden. In contrast to the Gulf of Aden, which is an elongated, narrow area with a simple bidirectional traffic flow, the Indian Ocean is much larger and crisscrossed, in all directions, by a multitude of traffic flows. This makes the design of an effective corridor system a complex optimization task.

Scenarios. We used the AgentC simulation to study three possible layouts of Indian Ocean corridor systems: (1) single *west-east corridor* channeling the large amount of west- and east-bound traffic (denoted as *Single-IO*), and (2) a more extensive *multi-corridor system* covering all the main traffic flows in the Indian Ocean (denoted as *Multi-IO*). Results are compared with the current setup where no corridors are used in the Indian Ocean (denoted as *None-IO*). IRTC corridor is considered in all cases. See Figure 4 for corridor layouts.

[9]http://agents.fel.cvut.cz/projects/agentc

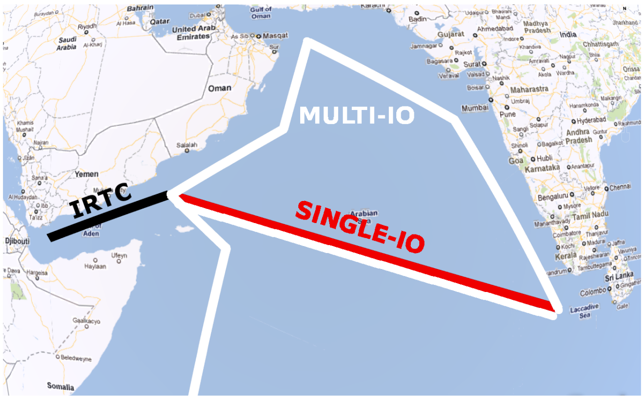

Figure 4: Corridor layouts for the Indian Ocean corridor system. The Single-IO layout uses the IRTC and the red east-west corridor only; the Multi-IO layout utilizes all depicted corridors.

In addition, we were interested in better understanding the possible synergy of deploying navy vessels alongside transit corridors. As a second study parameter, we therefore varied the number of deployed navy warships, using the stationary deployment method to determine their positions (see Section 4.3). All results presented are for three active pirate attack groups.

Results. We have evaluated all performance metrics defined in Section 2.4. Average transit distance and duration only depend on the corridor system and amounted to 6696km / 237h for no corridors in the Indian Ocean, 6703km / 237h for the Single-IO and 6819km / 242h for the Multi-IO corridor setup.

Pirate attacks statistics depend on both study variables. The numbers of hijacks in Figure 5 confirm the synergistic effect of transit corridors and patrolling—the Multi-IO corridor setup boosts protection force of patrols up to 40% in the case of 100 patrols (28.9 vs. 20.4 hijacks for None-IO and Multi-IO corridor system, respectively). A detailed breakdown of attack outcomes for the multi-IO corridor configuration (Figure 6) indicates that the decrease in hijacks is both due to the warship deterrence effect and the ability to intercept pirate attacks if they actually take place (more of the latter as the number of patrols increases). Finally, in Figure 7 we compare geographical distribution of vessel hijacks for None-IO and Multi-IO corridor setup with 20 patrolling warships. The distribution clearly depicts a high-risk hotspot north-east of the Socotra island.

Overall the results suggest that establishing a transit corridor system in the Indian Ocean is an effective way of increasing the security of transit on the expense of a very small increase in transit distance and duration (about 2% in the case of Multi-IO configuration). Further improvements might be attained if corridors are combined with group transit and/or escorted convoy measures, though that could have a noticeable impact on transit duration.

5.2 CS II: GOA Group Transit Optimization

In August 2010, the Group Transit Scheme was introduced to further reduce the risk of pirate attacks on vessels transiting the Gulf of Aden. The scheme, which is closely related to the IRTC, groups vessels traveling at similar speeds so

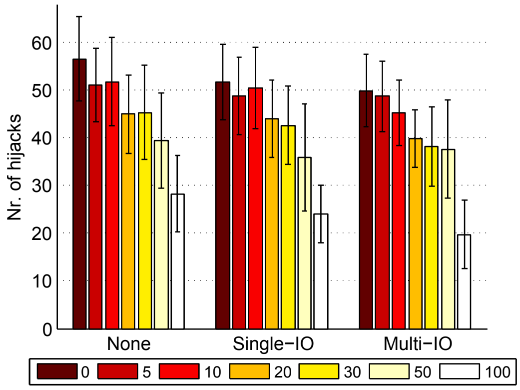

Figure 5: Dependency of the number of hijacks on the corridor system and the number of patrols (0–100). Standard deviation over 100 simulation runs also depicted.

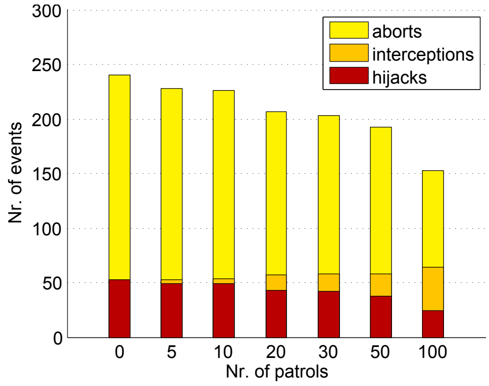

Figure 6: Ratios of *hijack*, *abort* and *interception* outcomes of pirate attacks for different numbers of patrol vessels.

that they cross high-risk areas close together as this provides additional deterrence to pirates and facilitates military response in case of an attack. Each transit group follows a recommended route through the IRTC at a published speed and fixed schedule (see Figure 8a) designed to maximize protection in highest risk times and areas. The schedule specifies vessel entry times depending on vessel's cruising speed. Five speed levels and consequently five speed groups are currently used—10, 12, 14, 16, and 18 knots.

The current number of speed levels and their uniform spacing is not optimum, given the distribution of cruising speeds in typical transit traffic (see Figure 8b). The aim of this study thus was to find out whether a different distribution of speed levels could reduce the transit delay incurred by following the group transit schedule.

Scenarios. The study variables were the number and distribution of speed levels used by the Gulf of Aden group transit scheme. For each number of speed levels, their optimum distribution was determined as the one maximizing the average speed of transit (and thus minimizing average transit time).

In addition to proposing optimum speed level distribution for fixed-schedule group transit, we have also explored the potential of a dynamic, negotiation-based group transit scheme in which groups are formed on-the-fly as vessels arrive.

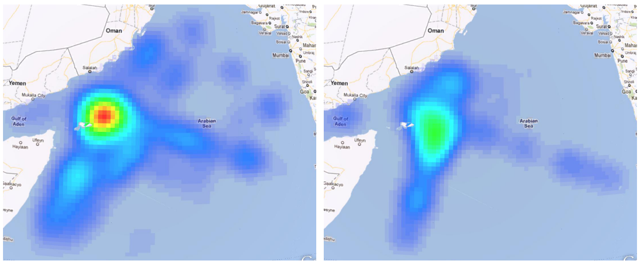

Figure 7: Geospatial distribution of hijacks for None-IO (left) and Multi-IO (right) corridor system configurations.

Results. We used the simulation to evaluate all key performance metrics. Surprisingly, no significant shift in the number of hijacks was observed. Upon a closer analysis, this can be explained by two forces cancelling each other—on one hand, a higher number of speed levels results in higher average transit speeds which reduces the attack success rate; on the other hand, more speed levels means more transit groups which makes their protection more difficult.

Expectedly, the average transit delay decreases with increasing the number of speed levels (Figure 8c). For the approximately 22 thousands vessels transiting the Gulf of Aden every year, the optimized 6 speed-level transit scheme would save over 400 days of total transit time a year. Considering an average daily vessel operational cost of US$30K, this translates into savings of approximately US$12mn a year. The dynamic group transit scheme surpasses even the best fixed-schedule scheme both in the number of hijacks and the transit time (the red dashed line in Figure 8c). Practical application of dynamic group transit would, however, require more extensive changes to the way transit is organized—a rather symptomatic trade-off between the performance and practicality of piracy counter-measures (see also the next section). See [17] for a more detailed evaluation and discussion of the case study.

6. LESSONS LEARNED

Overall, the multi-agent paradigm proved very useful during all stages of the development process—providing a conceptual framework for the analysis of the problem, informing architectural decisions during system design and, finally, supplying modeling and optimization techniques to implement the required functionality. Many challenges were not technical; often they lied in the (in)ability to obtain essential domain knowledge and datasets.

That said, the development of the simulation would have been easier if there was an agent-based simulation platform offering higher-level abstractions for representing individual- and collective level behavior and capable of simulating thousands of agents simultaneously. As mentioned in Section 3.2, to our best knowledge, no such a domain-independent platform currently exists, despite the fact that it would be useful in a wide range of applications. The optimization part would benefit from further research on (route) planning, scheduling and security resource allocation in hostile settings. Existing techniques, grounded largely in computational game theory, remain limited in their scalability and reliance on strong assumptions concerning rationality of adversaries and observability of their actions, although there is promising recent work on their relaxation (e.g., [7]).

As far as pitching of the agent-based approach is concerned, the ability to visualize the execution of individual simulation runs proved vital. First, it aided in conveying the very idea of maritime transportation as a multi-agent system. Second, it helped in winning the confidence of domain experts by allowing them to peek inside the model (see Figure 2) and verify that it behaves realistically on the micro-level. A key selling point of the agent-based approach to analyzing counter-piracy measures was the ability of the approach to analyze *hypothetical what-if* scenarios not yet occurring in the real world. Such scenarios cannot be reliably explored using standard statistical and/or data mining methods because they are too different from existing real-world situations on which datasets required for generating such models can only be obtained.

Working with the user community, we were constantly reminded of the necessity to maintain a proper balance between the quest for sophisticated, optimum solutions and their suitability for real-life deployment. Simple, suboptimal yet robust solutions can often be more suitable not only because they rely on fewer uncertain assumptions (such as rationality, observability or information sharing), but also because they are easier to explain and compatible with the existing infrastructure and processes in the generally very conservative maritime domain. Rather than coming up with a revolutionary optimum ways of managing counter-piracy measures, a more evolutionary approach seems more suitable, starting from concepts and measures already familiar to domain experts and using sophisticated techniques to discover their optimum configurations.

In contrast to other transportation domains, (global) maritime shipping seems underrepresented in the applied research on multi-agent systems and, in fact, on computational modeling and optimization in general. This is despite the fact that the global, transnational nature of maritime shipping and the consequent lack of strong central authorities, complex incentive structure and stiff economic competition make the multi-agent framework indispensable for accurately representing global shipping problems. Given the size of the shipping industry, estimated at several hundred US$ billion annually, current situation presents a sizable opportunity for innovative applications of multi-agent techniques.

7. RELATED WORK

Unlike in other areas of transportation, most notably road and air transportation, the deployment of computational modeling in the maritime domain is limited. Existing work either focuses on traffic in ports and national, coastal waters [8, 5] or uses high-level equation-based models [1] unfit for capturing individual-level behavior and inter-vessel interactions essential to model maritime piracy. The relative lack of work addressing global shipping as a whole is partly due to the global, international nature of maritime shipping and the consequent lack of a strong authority that would drive implementation of such methods.

Focusing on the very phenomenon of maritime piracy, the work is even more slim and concentrated primarily in the fields of security studies, international relations and global policy (e.g. [11]). Only very recently, initial attempts at applying computational modeling and optimization to maritime piracy have emerged but focus exclusively at military aspects of the problem [14, 12, 4].

Speed	Entry point A – time	Entry point B – time
10 kts	04:00 GMT+3	18:00 GMT+3
12 kts	08:30 GMT+3	00:01 GMT+3
14 kts	11:30 GMT+3	04:00 GMT+3
16 kts	14:00 GMT+3	08:30 GMT+3
18 kts	16:00 GMT+3	10:00 GMT+3

(a) Current group transit schedule.

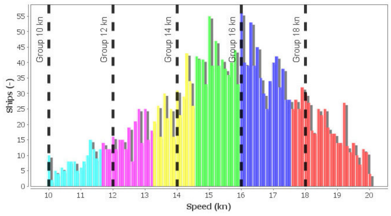

(b) Transit vessel speed distribution.

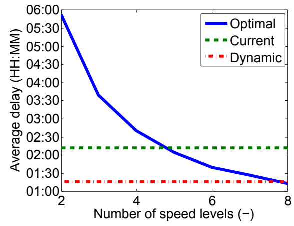

(c) Average transit delay.

Figure 8: (a) Schedule used by the current Gulf of Aden group transit scheme, (b) histogram of transit traffic speeds and its current (dashed lines) and optimum (colors) binning (for 6 speed levels), (c) reduction of the average transit delay with the increasing number of speed levels; delay for the current suboptimum (green) and dynamic grouping (red) schedule also shown.

8. CONCLUSIONS

We have shown how the multi-agent approach can be used to conceptualize and consequently address important challenges in planning and managing counter-piracy activities. A combination of multi-agent simulation and optimization proved very useful, enabling to evaluate and optimize the performance of counter-measures in a wide range of what-if scenarios. To our best knowledge, our work is the first integrated application of agent-based techniques to high-seas maritime security and, in fact, to global shipping analysis and optimization in general. The techniques developed enable commanders, policymakers and other relevant stakeholders to make better, more informed decisions and to improve maritime security with reasonable additional cost.

Acknowledgements

Funded by the Office of Naval Research (grant no. N000140 910537) and by the Czech Ministry of Education, Youth and Sports (grant no. LH11051).

9. REFERENCES

[1] S. Bourdon, Y. Gauthier, and J. Greiss. MATRICS: A maritime traffic simulation. Technical report, Defence R&D Canada, 2007.

[2] B. Bošanský, V. Lisý, M. Jakob, and M. Pěchouček. Computing time-dependent policies for patrolling games with mobile targets. In *10th International Conference on Autonomous Agents and Multiagent Systems*, 2011.

[3] A. Bowden, K. Hurlburt, E. Aloyo, C. Marts, and A. Lee. The economic costs of maritime piracy. Technical report, One Earth Future Foundation, 2010.

[4] J. Decraene, M. Anderson, and M. Low. Maritime counter-piracy study using agent-based simulations. In *2010 Spring Simulation Multiconference*, page 165, 2010.

[5] K. Hasegawa, K. Hata, M. Shioji, K. Niwa, S. Mori, and H. Fukuda. Maritime traffic simulation in congested waterways and its applications. In *4th Conference for New Ship and Marine Technology, Chine*, pages 195–199, 2004.

[6] O. Hrstka and O. Vaněk. Optimizing group transit in the Gulf of Aden. In *15th International Student Conference on Electrical Engineering*, 2011.

[7] D. Korzhyk, V. Conitzer, and R. Parr. Solving Stackelberg games with uncertain observability. In *10th International Conference on Autonomous Agents and Multiagent Systems*, 2011.

[8] E. Köse, E. Basar, E. Demirci, A. Güneroglu, and S. Erkebay. Simulation of marine traffic in Istanbul strait. *Simulation Modelling Practice and Theory*, 11(7-8):597–608, 2003.

[9] Y. Linde, A. Buzo, and R. Gray. An algorithm for vector quantizer design. *IEEE Transactions on Communications*, 28(1):84–95, 1980.

[10] P. Novák. Jazzyk: A programming language for hybrid agents with heterogeneous knowledge representations. *6th International Workshop on Programming Multi-Agent Systems*, pages 72–87, 2009.

[11] F. C. Onuoha. Piracy and maritime security off the Horn of Africa: Connections, causes, and concerns. *African Security*, 3(4):191–215, 2010.

[12] L. Slootmaker. Countering piracy with the next-generation piracy performance surface model (master thesis). Technical report, Naval Postgraduate School, Monterey California, 2011.

[13] F. J. Sluiman and H. de Konig. Naval vessel traffic services: Enhancing the safety of merchant shipping in maritime security operations. *Naval War College Review*, 63(3):123–137, 2011.

[14] T. Tsilis. Counter-piracy escort operations in the Gulf of Aden (master thesis). Technical report, Naval Postgraduate School, Monterey California, 2011.

[15] H. Van Dyke Parunak, R. Savit, and R. Riolo. Agent-based modeling vs. equation-based modeling: A case study and users' guide. In *1st International Workshop on Multi-Agent Systems and Agent-Based Simulation*, pages 277–283, 1998.

[16] O. Vaněk, B. Bošanský, M. Jakob, and M. Pěchouček. Transiting areas patrolled by a mobile adversary. In *2010 IEEE Symposium on Computational Intelligence and Games*, pages 9–16, 2010.

[17] O. Vaněk, M. Jakob, O. Hrstka, and M. Pěchouček. Using multi-agent simulation to improve the security of maritime transit. In *12th International Workshop on Multi-Agent-Based Simulation*, pages 12–23, 2011.

[18] O. Vaněk, M. Jakob, V. Lisý, B. Bošanský, and M. Pěchouček. Iterative game-theoretic route selection for hostile area transit and patrolling. In *10th International Conference on Autonomous Agents and Multiagent Systems*, pages 1273–1274, 2011.