Electronic supplement

To err is robotic, to tolerate immunological: fault detection in multirobot systems

Danesh Tarapore[1,2,3], Pedro U Lima[2], Jorge Carneiro[1], Anders Lyhne Christensen[4,5]

[1] Instituto Gulbenkian de Ciência, Rua da Quinta Grande 6, 2780-156 Oeiras, Portugal.

[2] Instituto de Sistemas e Robótica, Instituto Superior Técnico, Av. Rovisco Pais 1, 1049-001 Lisboa, Portugal.

[3] Institut des Systèmes Intelligents et de Robotique, Université Pierre et Marie Curie 6, CNRS UMR 7222, F-75252, Paris Cedex 05, France.

[4] Bio-inspired Computation and Intelligent Machines Lab, Avenida das Forças Armadas, 1649-026 Lisboa, Portugal.

[5] Instituto de Telecomunicações, Instituto Universitário de Lisboa (ISCTE-IUL), Avenida das Forças Armadas, 1649-026 Lisboa, Portugal.

E-mail: `daneshtarapore@gmail.com`

1. Effects of simplification to CRM implementation

1.1. Performance

An analysis of possible detrimental effects to fault-detection performance, because of simplifications to the CRM implementation (see Section 3), was conducted by performing the following four experiments with the *complete CRM implementation* (detailed in Section 2): (a) detection of fault-simulating behaviours (Figure S1); (b) tolerance to normal behaviours (Figure S2); (c) tolerance to transitions in normal behaviour of entire swarms (Figure S3); and (d) response (immune and tolerance) to transitions in single agent behaviour (see Table 5 for details of all four experimental setups and the quantification of performance).

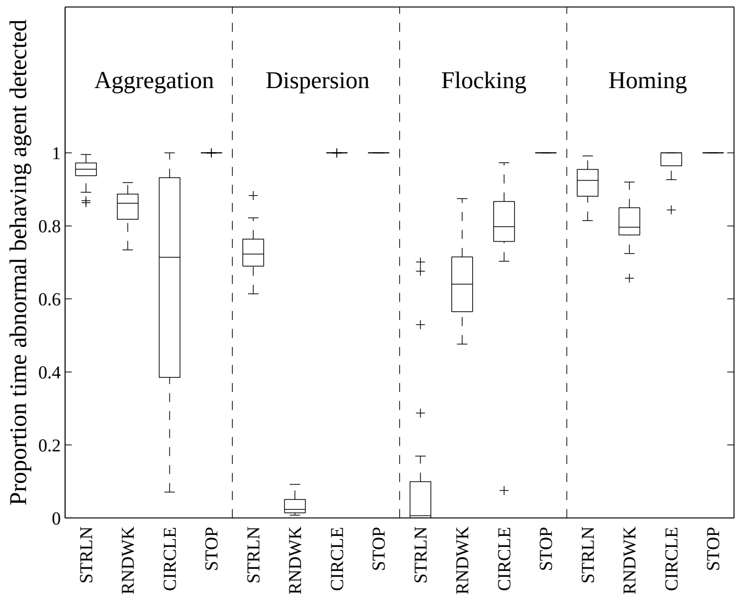

Figure S1: **Detection of abnormalities, when employing the complete CRM implementation.** Proportion of time fault-simulating agent is detected as abnormal across 20 replicates, in each of the 16 distinct combinations of normal (aggregation, dispersion, flocking and homing) and abnormal (STRLN, RNDWK, CIRCLE and STOP) behaviours.

In the first three experiments, the difference in performance between the simplified (Figure S1, S2 and S3) and the complete CRM implementation (Figure 4, 5 and 6) was not different from 0 (two-sided sign test, $d.f. = 19$, all $p > 0.47$). For the fourth experiment, the performance during the first half of the simulation was 0.99 ± 0.001 (mean±SD), in the two CRM implementations. Furthermore, in the second half of the simulation, the performance was identical at 0.70 ± 0.386 in both implementations. In summary, the simplifications to the implementation of the CRM did not have any effect on the performance of the model in the detection of agents executing fault-simulating behaviour, and the tolerance of agents behaving normally.

1.2. Computational cost

The computational resources required for running the simplified and the complete CRM implementations was compared by recording the number of floating-point operations executed by the two CRM-based abnormality

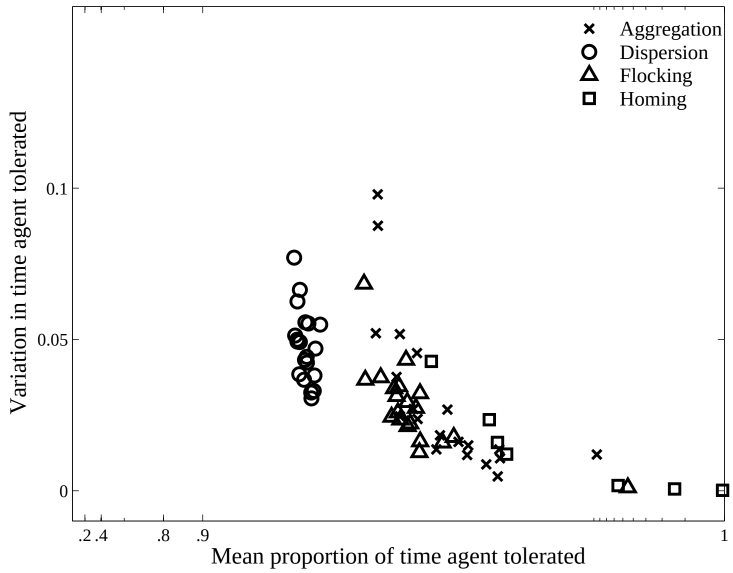

Figure S2: **Tolerance to normal behaviour with the complete CRM implementation.** Mean and variation in proportion of time agent tolerated, across the 20 agents of the MAS, in each of 20 replicates of four swarm behaviours. Variation measured in each replicate as the difference between the maximum and minimum time tolerated, across the 20 agents.

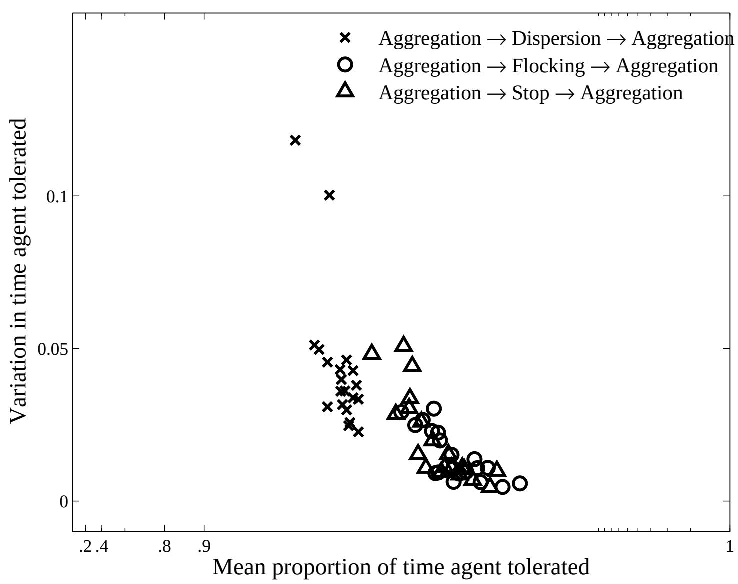

Figure S3: **Tolerance to changes in normal behaviour with the complete CRM implementation.** Mean and variation in proportion of time agent tolerated, across the 20 agents of the MAS, in each of 20 replicates, and two transitions in normal behaviour: (a) aggregation to dispersion to aggregation (crosses), (b) aggregation to flocking to aggregation (circles), and (c) aggregation to STOP to aggregation (asterisks).

detectors. The Figure S4 shows the number of floating-point operations (FLO) in experiments involving the transition of normal behaviour from aggregation to dispersion to aggregation.

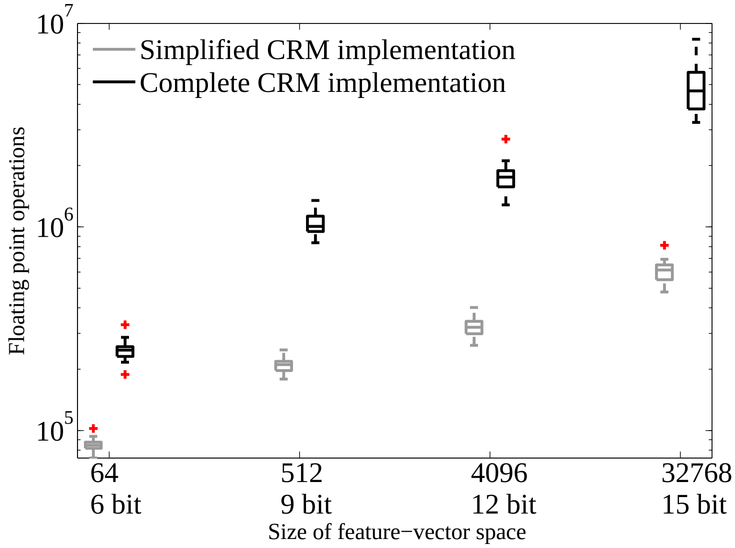

Figure S4: **Scalability of resources for the simplified and complete CRM implementations in terms of FV classification space.** Computational costs of abnormality detection with the two CRM implementations, for FV space of 6, 9, 12 and 15-bits FV length (in log scale), across 20 replicates. Each box corresponds to the average number of FLOs per agent, per control cycle in experiments involving the transition of normal behaviour from aggregation to dispersion to aggregation (details of experimental setups in Table 5).

In both the simplified and the complete CRM implementation, while the size of the FV space increases exponentially with the number of features considered, the computational resources in terms of FLOs required to execute the CRM-based abnormality detection increases in accordance with a power function. The exponent of the power function was 0.3 ($R^2 = 0.97$) for the simplified, and 0.4 ($R^2 = 0.99$) for the complete model implementation.

2. Detailed comparison between the CRM and the threshold-based model

2.1. Performance

Across all the four experimental setups, and the FVs of length 6, 9, 12 and 15-bits, abnormality detection with the simplified CRM implementation outperformed the threshold-based model (see Table S1). In all but one of these 16 treatments, the difference in performance was significant (Mann-Whitney-Wilcoxon nonparametric test, all $p < 0.001$), with the estimated effect size exceeding 0.79 (out of 1). In the experiment involving the change in behaviour of a single agent (setup D, see Table 5), our CRM implementation still outperformed the threshold-based model, but the difference was no longer significant ($p = 0.54$). The effect size is computed as the ratio of the Mann-Whitney U statistic and the product of the two sample sizes (20×20). This statistic estimates the probability that the performance of a replicate randomly drawn of the CRM outperforms a replicate randomly drawn of the threshold-based model.

Table S1: Performance of the simplified CRM implementation and the threshold-based model.

Exp. setup	FV length	Performance (Mean±SD)		p-Value*	Effect size**
		CRM implementation	Threshold-based model		
A	6 bits	0.90 ± 0.03	0.85 ± 0.02	1.9×10^{-5}	0.88
	9 bits	0.97 ± 0.02	0.88 ± 0.02	3.9×10^{-8}	1.0
	12 bits	0.94 ± 0.02	0.88 ± 0.02	2.3×10^{-7}	0.97
	15 bits	0.94 ± 0.02	0.90 ± 0.02	6.3×10^{-6}	0.91
B	6 bits	0.99 ± 0.001	0.96 ± 0.01	3.4×10^{-8}	1.0
	9 bits	0.97 ± 0.001	0.86 ± 0.01	3.4×10^{-8}	1.0
	12 bits	0.94 ± 0.001	0.78 ± 0.02	3.4×10^{-8}	1.0
	15 bits	0.94 ± 0.001	0.63 ± 0.02	3.4×10^{-8}	1.0
C	6 bits	1.0 ± 0.001	0.96 ± 0.01	3.4×10^{-8}	1.0
	9 bits	0.99 ± 0.002	0.88 ± 0.01	3.4×10^{-8}	1.0
	12 bits	0.99 ± 0.002	0.76 ± 0.02	3.4×10^{-8}	1.0
	15 bits	0.99 ± 0.004	0.59 ± 0.04	3.4×10^{-8}	1.0
D	6 bits	0.85 ± 0.2	0.80 ± 0.2	0.33	0.54
	9 bits	0.85 ± 0.2	0.76 ± 0.2	8.1×10^{-4}	0.79
	12 bits	0.86 ± 0.2	0.69 ± 0.15	8.1×10^{-4}	0.8
	15 bits	0.86 ± 0.2	0.59 ± 0.1	8.1×10^{-4}	0.83

*Significance computed with the Mann-Whitney-Wilcoxon nonparametric test, performed at the 0.05 significance level.
**Effect size is estimated as the ratio of the Mann-Whitney U statistic and the product of the two sample sizes.

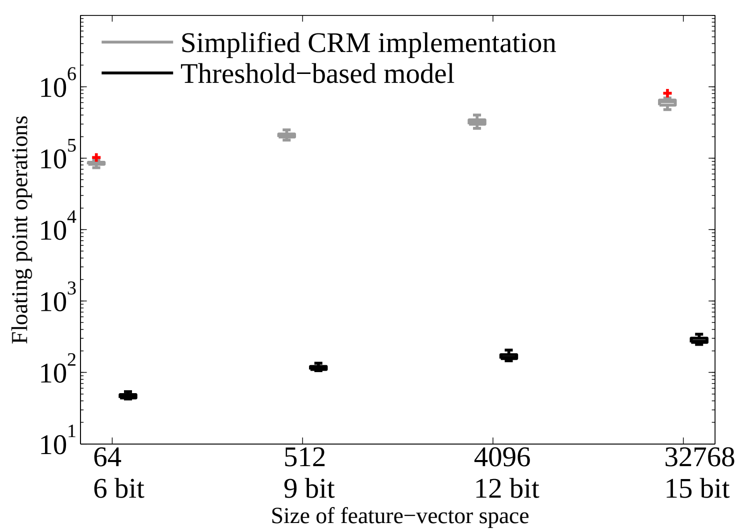

Figure S5: **Scalability of resources for the simplified CRM implementation and the threshold-based model, in terms of FV classification space.** Computational costs of abnormality detection with the CRM and the threshold-based model, for FV space of 6, 9, 12 and 15-bits FV length (in log scale), across 20 replicates. Each box corresponds to the average number of FLOs per agent, per control cycle in experiments involving the transition of normal behaviour from aggregation to dispersion to aggregation (details of experimental setups in Table 5).

2.2. Computational cost

Despite the poor performance in fault-detection with the threshold-based model, the computational resources required for its execution were minimal in comparison to the CRM. In the experiments involving the transition of normal behaviour from aggregation to dispersion to aggregation, the number of FLOs required to execute the threshold-based model was at least two orders of magnitude lower than the CRM, irrespective of the size of the FV classification space (see Figure S5). Furthermore, the exponent of the power function describing the increase in the number of FLOs with increments in the size of the FV space was 0.25 ($R^2 = 0.98$) for the the threshold-based model (exponent at 0.3 for simplified CRM implementation).