Managing Renewable Resources Via Collective Intelligence

Brede, M., F. Boschetti, and D. McDonald

CSIRO - Marine and Atmospheric Research, Australia
Email: [Markus.Brede,Fabio.Boschetti,David.Mcdonald]@csiro.au

Keywords: ABM, Game Theory, Resource Exploitation, Optimisation.

EXTENDED ABSTRACT

In a recent set of work (Boschetti, 2007; Brede et al, 2007; Brede and De Vries, 2007) we have explored the potential for porting tools originally developed for the optimisation of multi-agents engineering systems (Wolpert and Tumer, 2001; Wolpert et al, 2004) into resource management modelling. These tools were designed to minimise interference between system components which may arise in system optimisation, as can be found in data communication problems and aerospace engineering (Bieniawski et al, 2004; Wolpert et al, 2000; Wolpert and Tumer, 1999).

The analogy with resource exploitation and management is based on viewing human agents (fishers, farmers, etc..) as members of a larger system (fleet, farming community, etc,) which accesses a common resource. The global harvest is the sum of the harvest of each agent. However, maximising the harvest of each agent does not necessarily result in a maximum global harvest (Arthur, 1994; Challet and Zhang, 1998). Agents naturally tend to compete for a better resource allocation; when the resource is spread over several zones, resource availability at specific locations may attract more agents than necessary for the desirable level of exploitation of the zones' resource, thereby worsening the global harvest. Also, overuse of the resource by a few agents will result in the collapse of the entire resource and, consequently, poorer performance for each agent. Clearly dynamics of various levels are connected: the action of each agent impinges on the performance of the entire community and, vice-versa, the global performance of the entire community will affect future returns for each individual (Batten, 2007).

It is easy to see (Boschetti, 2007) that, given sufficient resources, optimal global exploitation can be achieved by spreading the community harvesting effort proportionally to the resource availability in different zones. However, achieving this optimal allocation without centralised control is not trivial and here is where ideas inherited from engineering, computer science and game theory can be useful.

Of course, human agents are very different from components of an engineering system: their decision making is not mechanical and thus can not be written down as an algorithm. Often decision making is not fully rational, that is, it is not based solely on economic incentives. It follows that not all tools which perform well in engineering problems may be easily ported to human applications. The tool we discuss in this paper, the Collective Intelligence (COIN; Wolpert et al, 2000), gives some confidence that it might be an exception: while it cannot overcome the 'non algorithmic' component of human decision making, its functioning is so simple that it could be implemented by pen and paper via very simple bookkeeping and accounting (Boschetti, 2007), making it easily accessible to human agents without any need for computers or mechanical aids.

Initial numerical tests have proved encouraging in different scenarios of resource dynamics and exploitation pattern (Boschetti, 2007; Brede et al, 2007; Brede and De Vries, 2007). Our on-going research aims to explore within which context this technique can be used as a resource management tool. As a step in addressing this problem, in the present paper we envisage a large fishing fleet targeting a resource spread over a number of fishing zones. Given a certain management goal, in this case optimal resource exploitation, we ask whether a mixed centralised resource allocation, in which fleet managers adopt COIN to direct fishing vessels towards chosen fishing zones, can do better than a totally decentralised approach in which each vessel adopts COIN for its own decision making. Both approaches are compared against what we define as the 'null hypothesis' approach, in which both individual vessels and fleet managers choose to optimise 'greedily' a perceived immediate profit.

1. THE RESOURCE EXPLOITATION MODEL

In this paper we report a model of a simplified, non-spatially explicit fishery. We imagine N fishing vessels, belonging to $2 \leq M \leq N$ fleet managers: managers can own either a single vessel or a larger fleet. We also imagine Z fishing zones in which an amount $Fish_{z}, z = 1..Z$ of resource is available. The managers do not have information about the global distribution of $Fish_z$ and decide where to direct their vessels according to the discounted returns of past catches in the different fishing zones (see Boschetti, 2007, for details).

At each fishing period a vessel targets a single fishing zone. We assume that the decision on which zone to target is taken solely by the vessel's manager. Once the zone has been targeted, the resource available in that zone is shared equally among all vessels accessing it. Each vessel has a maximum fishing capacity (which can also be interpreted as a maximum allowed quota). Thus, the catch of a vessel n is given by

$$Catch_n = Min(Fish_{Z_n} / Crowd_{Z_n}, Quota) \qquad (1)$$

where $Catch_n$ is the amount of fish caught by vessel n, Z_n is the fishing zone chosen by vessel n, $Fish_{Z_n}$ is the amount of fish available in $zone_n$, $Crowd_{Z_n}$ is the number of vessels $\qquad (1)$

$\qquad (1) \qquad (1)$which chose to fish in Z_n, with which vessel n has to share the available resource. We do not model fishing costs (navigating to the zones, equipment renting/buying, etc) though these could be included easily if needed.

The total catch of the fleet is obviously given by the sum of each vessel's catch,

$$TotalCatch = \sum_{n=1,N} Catch_n \qquad (2)$$

The maximum possible catch of the entire fleet is given by either the total amount of resource in the fishery or by the sum of the maximum allowed quota per vessel, if the resource is abundant:

$$MaxCatch_{Fleet} = Min(\sum_{z=1,Z} Fish_z, N*Quota) \qquad (3)$$

This is the optimal catch against which the results of the fishing strategies will be evaluated. Notice that, because each vessel has a maximum allowed quota, we have:

$$TotalCatch \leq MaxCatch_{Fleet} \qquad (4)$$

that is, unless the vessels spread their effort wisely, the fleet may not be able to catch to its full capacity.

2. THE COLLECTIVE INTELLIGENCE

As explained above, the exploitation of the resource described in the previous section can be seen as an optimisation problem, in which we aim to allocate the N vessels proportionally to the resource $Fish_z$.

From an optimisation perspective, we could choose to optimise two quantities: a) the 'private' return for each vessel and b) the 'global' return to the entire fleet. Optimising either of them in isolation is known to be sub-optimal. The 'private' return is optimised in a class of problems known as a Minority Game (Challet and Zhang, 1998; Arthur, 1994), in which it is shown that the fleet never reaches a distribution proportional to the resource distribution; rather it oscillates around the optimum value and these oscillations correspond to a waste of resource (Boschetti, 2007) and never dissipate. Alternatively, the 'global' return is optimised in what is called a 'team game' and it is known that it reaches an optimal exploitation distribution only for very small problems (Wolpert and Tumer, 2001).

The COIN approach can be seen as a sort of compromise between these two. It is based on each vessels aiming to optimise their own 'private' return, but in this case the private return is a function of how they influence the 'global' return. In particular, each vessel tries to maximise its *impact* on the global return, where the impact is equal to the difference between the global catch for the fleet and the catch that the fleet *would have caught* had the vessel not being present. The presence of a maximum fishing capacity (or a quota) for each vessel may result in this impact be different from the catch of the vessel itself. Details on how this quantity can be calculated are given in the Appendix and we refer the reader to Boschetti (2007) for more details. Importantly, this measure can be calculated using only local information about the area targeted by the vessel, without any need for global information.

3. PREVIOUS RESULTS

In this section we briefly summarise the results we obtained in previous tests of COIN performance in fisheries resource exploitation problems and we refer the readers to the original publications for further details.

In Boschetti (2007) we describe a number of virtual experiments mimicking a fishing fleet operating in areas of different fishing capacity, but fully renewable resource. In that scenario, COIN provides optimal catches for the fleet while, at the

same time, each individual vessel also maximizes its own profit: in principle COIN requires fishing vessels not to act greedily but no individual sacrifice is required to achieve the common goal. We also show that a fleet following a COIN strategy adapts much faster to change in resource distributions, promising increased benefits over standard approaches in volatile environments.

These results were extended to scenarios in which the resource dynamics were explicitly modelled (Brede et al, 2007). We also explored the trade-off between long and short term planning by providing vessels with some knowledge of the time evolution of the resource and thereby allowing them to plan their fishing behaviour in light of predicted long term resource behaviour (Brede and de Vries, 2007). Finally, vessels were allowed to dynamically choose what strategy to adopt (ranging from fully cooperative, fully competitive, random and COIN) according to the strategy past return, in a typical evolutionary economics scenario (Gintis, 2000). We showed that the balance of vessels choosing a COIN versus a greedy approach depends crucially on the resource availability. Importantly, we also showed that although COIN requires more local information than a fully greedy approach, COIN is not very sensitive to incorrect/false information, providing acceptable performances even in the presence of considerable noise (Brede and de Vries, 2007). This reinforced our optimism that the method could, in principle, be implemented in real scenarios.
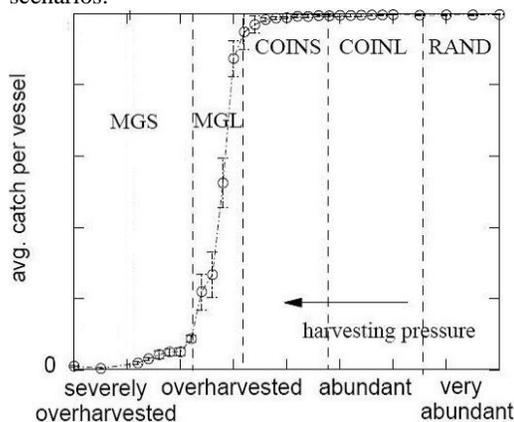

Figure 1. Average catch per vessel (Y) as a function of resource availably (X). The vertical dashed bars indicate which strategy is optimal for different resource abundance; RAND= random choice of fishing zone, COINL= COIN with long term projection of resource dynamics, COINS=COIN without projection of resource dynamics, MGL= greedy strategy with long term projection of resource dynamics, MGS= greedy strategy without projection of resource dynamics.

The dependence of COIN performance on the resource availability is of particular interest for 3 reasons: a) it indicates when COIN can be most effective, b) it can give an indication of resource availability (Brede et al, 2007) and c) it can describe the dynamics of cooperation versus competition in natural resource problems. These ideas are summarised in Figure 1 (see Brede et al, 2007 for more details).

The X axis gives an indication of resource abundance in relation to the fishing capacity of the overall fleet, while the Y axis shows the modelled average catch per vessel (the catch plateaus for abundant resources because of the limit in vessel fishing capacity due to physical limitations or quotas). The vertical lines indicate which fishing strategies give the best catches under different resource regimes. Starting from the right-hand side and moving leftward, we can see how the optimal fishing strategy changes as a function of resource abundance. When the resource is very abundant (basically unlimited in comparison to the fleet capacity) there is no need to put much effort in choosing where and how to fish, and consequently a random strategy (RAND in the figure) performs well. When the resource is abundant, but not unlimited, a COIN strategy accounting for long term resource dynamics (COINL) is best. From now on, further resource reduction favours more and more greedy fishing strategies; when the resource is limited but not over-exploited COIN with no long term projection of resource dynamics fares best (COINS), while when the resource gets overexploited fully competitive, greedy behaviours become optimal and thus more dominant in the fleet (MGL, greedy behaviour with long term projection and MGS, greedy behaviour with no long term projection). Important for our discussion is the transition from COIN to greedy strategies, which maps the transition between under-exploited and over-exploited resources.

Suppose that this transition could actually be detected in practise and consequently be used as an indicator of the resource state. If each vessel is managed independently, this knowledge would not change the decision of whether or not to go fishing and consequently potentially inflict long-term damage to the resource. Effectively, an over-exploited resource results in a Tragedy of the Commons situation (Hardin, 1968), in which little economic benefit is provided to avoid over-exploiting a limited resource. However, imagine vessels are managed as members of larger fleets; then it may make economic sense to prevent certain vessels from going fishing, thus reducing the cost of fishing by limiting the size of the fleet to the minimum required to catch the limited

resource. This would not reduce the environmental damage but would reduce the cost of exploitation.

It is obvious from Section 2 that the impact of a fishing vessel, as defined by the COIN approach, is exactly the piece of information which a fleet manager needs in order to choose the optimal fleet size. Before addressing this problem though, we need to make sure that the COIN approach can be used by a fleet manager and that its performance is acceptable against a fully decentralised COIN and against a traditional greedy approach. This is the motivation of the tests we describe in the next section.
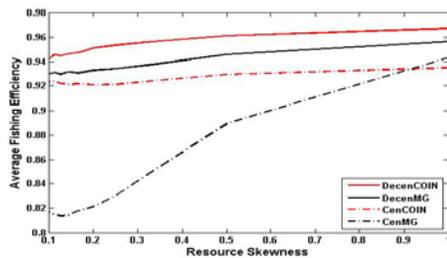

Figure 2. Fishing efficiency versus resource skewness for the 4 strategies under analysis.

4. FLEET MANAGEMENT BY COIN

The purpose of our tests is to compare 4 fishing strategies:

1) a fully decentralised COIN; in this case we have $M=N$, that is, each vessel is privately owned; at each period, each vessel decides where it will fish next according to the discounted record of past COIN *impacts.* We call this DecenCOIN in the coming figures.
2) A centralised COIN; the entire fleet is subdivided into smaller fleets, each managed separately. In these tests a sub-fleet comprises 10 vessels. At each period, the sub-fleet manager decides how to allocate his 10 vessels according to the discounted record of the past COIN *impacts* of his fleet. We call this CenCOIN.
3) A fully decentralised greedy fleet ($M=N$, each vessel is privately own); at each period, each vessel decides where to fish next according to the discounted record of past *catches.* We call this DecenMG (we use MG for consistency with the notation in our previous papers).
4) A centralised greedy fleet (each sub-fleet comprises 10 vessels). At each period, the sub-fleet manager decides how to allocate his 10 vessels according to the discounted record of the past *catches* of his fleet. We call this CenMG.

We want to check the relative performance of these strategies under different conditions of resource distribution, resource availability and fleet size.

4.1. Performance versus resource distribution and abundance

In Figure 2 we show the performances of the 4 scenarios described in the previous section for resource distribution of varying skewness. All results are obtained by modelling a fleet of fixed size including $N=50$ vessels, which corresponds to 5 sub-fleets of 10 vessels for the centralised scenarios. In all figures the results are given for runs of 100 fishing periods, and include the initial transient during which the strategies train themselves on initially random data. We also model $Z=2$ fishing zones for which the resource ratio $Fish_2 / Fish_1 = k$, $k = 0.1, 0.2....1$; $k=1$ corresponds to an even distribution of resource while $k=1/10$ corresponds to a most skewed one. In the Y axis we plot the average fishing efficiency, that is the ratio between the average and the maximum possible catch per vessel. In the figure, red lines show the COIN results while black lines show the results of the greedy strategies. Also, thick lines show decentralised approaches and dashed lines show those that are centralised.

Figure 2 suggests a number of conclusions: first, decentralised strategies perform better than centralised ones, both for COIN and the greedy approach. Second, for both centralised and decentralised scenarios, COIN outperforms the greedy approach. Finally, the 4 strategies react differently to the different level of skewness in resource distribution. The centralised MG is strongly affected by it, the decentralised COIN and MG are only slightly affected and the centralised COIN almost unaffected.
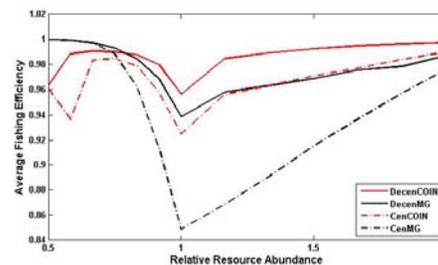

Figure 3. Fishing efficiency versus relative resource abundance for the 4 scenarios under analysis.

In Figure 3, we analyse the efficiency of the 4 scenarios versus relative resource abundance. This

is calculated as the ratio between available resource and the maximum fishing capacity of the fleet, ranging from *.5* (resource equal to half the fishing capacity) to *2* (resource equal twice the fishing capacity). In this case we modelled a moderately skewed resource distribution, $Fish_2 / Fish_1 = 1/3$. The behaviour of Figure 3 is more complicated than the one in Figure 2, showing that the resource abundance has a stronger influence on the fishing efficiency than its distribution. For a very scarce resource the greedy approaches outperform COIN, while the opposite is true for abundant resources, confirming the results displayed in Figure 1. This is true for both centralised and decentralised approaches. When the amount of resource is roughly equal to the fishing capacity, we have a minimum in fishing efficiency for all scenarios; this is the situation for which a proper fleet distribution is most crucial for optimal resource exploitation. Around this value, decentralised scenarios outperform centralised ones, as in Figure 2.

Finally, in Figure 4 we show a 2D plot in which the Y axis shows different levels of resource skewness and the X axis different levels of resource abundance. At each location on the plot we map which of the 4 scenarios performs best. In the figure dark tones refer to COIN and lighter tones to greedy strategies. As expected, for low relative resource abundances (<0.8 in the figure) greedy strategies perform better, with the centralised greedy strategy being optimal for more skewed resource distributions. For more abundant resources the COIN strategies perform better. In this case the decentralised COIN (black area) performs best in most scenarios.
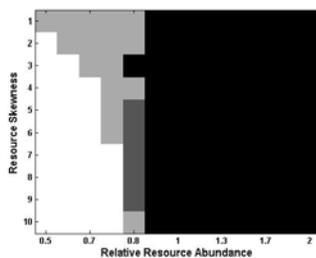

Figure 4. Best performing strategy as a function of resource skewness and relative resource Abundance. White=CenMG, light-gray=Decen_MG, dark-grey=CenCOIN, black=DecenCOIN .

4.2. Performance versus fleet size

Here we analyse how the results presented in the previous section scale with problem size. In Figure

5 we plot the average fishing efficiency versus relative resource abundance (as in Figure 3) for a much larger fishery, including a fleet of 200 vessels targeting a skewed resource distributed over 20 fishing zones.
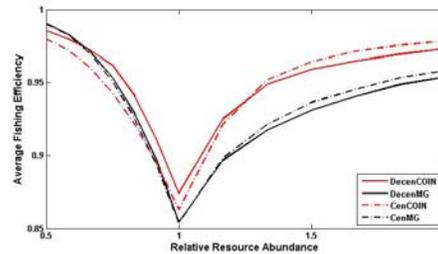

Figure 5. Fishing efficiency versus relative resource abundance for the 4 scenarios under analysis for a large fishery model.

The main difference between Figure 5 and Figure 3 lies in the better performance of the COIN strategies versus the greedy ones. For the larger fleet, COIN outperforms greedy strategies already for scarce resources, while for smaller fleets (Figure 3) this happens only for abundance resources. This suggests that the considerable increase in the number of options for the vessel allocation in the larger problem is handled better by the COIN than by the greedy approach. Similar results are reported in the COIN literature for engineering applications (Wolpert et al, 2004).
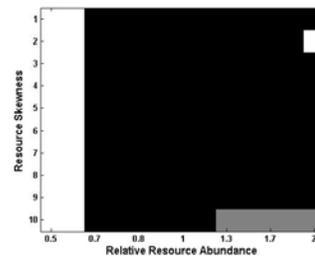

Figure 6. Best performing strategy as a function of resource skewness and relative resource Abundance for a large fleet.

Finally, in Figure 6 we show which strategies perform best for different values of resource skewness and resource abundance for the large fleet; this figure should be compared with Figure 4. Clearly, for larger problems, the decentralised COIN improves its role as dominant strategy.

5. DISCUSSION AND FUTURE WORK

The results presented in this paper can be summarised as follows:

1) except for very scarce resources, COIN outperforms greedy strategies; this is true both for centralised and decentralised approaches.

2) COIN performance scales up better than greedy approaches, and consequently promises to be particularly useful in the management of large problems.

3) Resource skewness does not seem to affect the strategies to a large extent, with the exception of a centralised greedy approach.

4) Most important for the scope of this work, decentralised strategies fare better than centralised strategies.

The last point deserves some discussion. At first sight, it suggests that a bottom-up approach is more effective than a top-down one. This would not prevent a manager adopting COIN to allocate fishing effort. It simply means that this manager should perform a COIN calculation at the scale of each single vessel belonging to its fleet rather than at the scale of its overall fleet.

There are a number of reasons behind the apparent contradiction in suggesting that the decentralised COIN approach could be implemented by a manager; that is in a pseudo-centralised fashion, rather than in a proper decentralised way, by each vessel's skipper. First, as mentioned in Section 3, we believe that the use of COIN by a management body would ease the potential introduction of the tool in resource management. The other reasons require a deeper investigation of the causes of the results presented above.

First, let's analyse the difference in performance between COIN and greedy approaches. As described in Section 2, the difference between the COIN and a greedy approach is that while the greedy algorithm accounts for each vessel's *catch*, the COIN accounts for the vessel's *impact* on the fleet overall catch. This allows COIN to decrease the number of vessels targeting areas which are already over-fished: in these zones the catch of a potential additional vessel would result in a decrease of the catch of other vessels targeting the same zone (since the zone is over-fished) which is why the vessel's impact on this zone would be zero (see Boschetti, 2007). The reason why COIN performs poorly against a greedy strategy in situations of scarce resources is that in these cases vessels are unable to produce an impact anywhere, since all zones are over-fished. As a result COIN lacks the information necessary to perform an appropriate effort allocation.

Now, let's direct our attention to the difference in performance between centralised and decentralised approaches. In a centralised approach, a fleet manager receives information about the potential impact of a vessel in each of the zones one of its vessels has targeted. In principle, the fleet manager has more information than each single skipper in his fleet. However, he does not have information on how to split the effort among the zones. This results in spreading the fleet equally over the zones which results in poorer performance.

In our research, however, we have assumed that each vessel will go fishing in each period. Let's consider the situation of a scarce resource in which a vessel forecasts it cannot make any impact in any fishing zone. As explained above, this results in lack of information to the COIN. However, in principle, this is a very useful piece of information telling the vessels that, for the good of the overall fleet, it should not go fishing: independently of where it will go fishing, the catch of the overall fleet will be unaffected. Obviously, if the vessel acts independently, in a bottom-up approach, there is no incentive for the vessel not to go fishing, because this would result in lost income. However, if the vessel is part of a larger fleet, there would be an incentive for the fleet manager to prevent the vessel going fishing, since his fleet would obtain the same overall catch with reduced costs.

The conjecture we suggest from the above discussion is that a fleet manager's optimal strategy would be to adopt a *decentralised* COIN in case of an abundant resource (which would guarantee optimal catches) and a *centralised* COIN in case of a scarce resource, thanks to which it could estimate the minimum number of vessels needed for the job, thereby minimising costs without impairing the catch itself. In order to carry out the above strategy, what is needed is a way to use the COIN impact information in order to decide which vessels should not be employed in a fishing period. We will direct our future work at exploring this idea and we hope to report on this at the conference presentation.

6. CONCLUSIONS

The relative performance of greedy approaches and COIN depend on resource availability and size of the problem, which suggests that an adaptive meta-strategy would probably be optimal in general resource management problems. A bottom-up decentralised strategy, in which calculation for optimal vessel allocation are carried out at the vessel level, rather than at the fleet level, provides best resource exploitation for most modelled scenarios.

Finally, we discussed how such a bottom-up approach could be implemented by a fleet manager, and conjectured that this may result in reduce exploitation costs.

7. REFERENCES

Arthur, W.B., (1994), Inductive behaviour and bounded rationality, *The American Economic Review*, 84, 406-411.

Boschetti, F., (2007), Improving resource exploitation via Collective Intelligence by assessing agents' impact on the community outcome, *Ecological Economics*, 63, 553-562.

Batten, D., (2007), Are some human ecosystems self-defeating?, *Environmental Modelling and Software*, 22(5), 649-655

Brede, M., F. Boschetti & D. McDonald, (2007), Strategies for Resource Exploitation, , *Ecological Complexity*, accepted.

Brede, M., & De Vries, (2007), Harvesting Heterogeneous Renewable Resources: Uncoordinated, Selfish, Team-, and Community-Oriented Strategies, *Environmental Modelling & Software*, in review.

Challet, D. and Y-C, Zhang, (1998), On the minority game: analytical and numerical studies. *Physica A*, 256, 514.

Gintis, H., (2000), Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction, Princeton University Press, Princeton, NY

(2000).Hardin, G., (1968), The tragedy of the commons, *Science*, 162, 1243-8.

Wolpert, D. and K. Tumer, (2001), Optimal Payoff Functions for Members of Collectives, *Advances in Complex Systems*, 4(2/3):265–279.

Wolpert, D., K. Wheeler, and K. Tumer, (2000), Collective Intelligence for Control of Distributed Dynamical Systems. *Europhys. Lett.*, 49 (6), p. 708.

Wolpert, D. and K. Tumer, (1999), An Introduction To Collective Intelligence, Tech Report: NASA-ARC-IC-99-63, http://ic.arc.nasa.gov/ic/people/kagan/coin pubs.html.

Wolpert, D., K. Tumer, and D. Bandari, (2004), Improving Search Algorithms by Using Intelligent Coordinates, *Physical Review E*, 69, 2004.

Bieniawski, S., D. Wolpert, & I. Kroo, (2004), Discrete, Continuous, and Constrained Optimization Using Collectives, AIAA Paper 2004-4580, *Presented at the 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Albany, NY, August 30-September 1, 2004.

APPENDIX A – COIN Impact Calculation

We define the impact of vessel *n1* on the overall fleet as the difference between the global catch of the fleet and the catch of the fleet *would have caught* if vessel *n1* had not gone fishing:

$$IMpact_{n1} = TotalCatch - TotalCatch^{-n1} \,, \qquad (A1)$$

where the superscript '-*n*' refers to the fleet without vessel *n1*. Notice that, because the catch of each vessel is limited by physical constraint or quota restrictions, we can have:

$$TotalCatch^{-n1} \neq TotalCatch - Catch_{n1} \cdot$$

Eq. A1 can be approximated by removing vessel *n1* from the fleet, leaving everything else unchanged. Let's suppose vessel *n1* targeted fishing zone *z1*. Clearly it could not have any impact on the other zones, so we need to concern ourselves only about zone *z1*. It thus follows that:

$$IMpact_{n1} = TotalCatch - TotalCatch^{-n1} = Fleet_{z1} * Min(Fish_{z1} / Fleet_{z1}, Quota) - Fleet_{z1}^{-n1} * Min(Fish_{z1} / Fleet_{z1}^{-n1}, Quota)$$

where *Fleet* is the size of the fleet which targeted zone *z1* and the possible catch per vessel is constrained by the quota. Since $Fleet_{z1}^{-n1} = Fleet_{z1} - 1$, we have thus have:

$$IMpact_{n1} = Fleet_{z1} * Min(Fish_{z1} / Fleet_{z1}, Quota) - (Fleet_{z1} - 1) * Min(Fish_{z1} / (Fleet_{z1} - 1), Quota)$$

which is the equation used in our calculation.

Clearly, for each fishing period, a vessel can obtain information only about the zone it has targeted. Information about the other fishing zones is available only via past catches. Each vessel stores fishing results in a table, in which catches or impacts from past periods are discounted linearly according to their age. The sum of the discounted catches (or impacts) for each zone, properly normalised, gives the probability of a vessel targeting that zone at next period. In order to prevent certain zones never being targeted by a certain vessel, probabilities below a given threshold are forbidden.