An Introduction to Swarm Intelligence Issues

Gianni Di Caro

gianni@idsia.ch

IDSIA, USI/SUPSI, Lugano (CH)

Topics that will be discussed

- Basic ideas behind the notion of Swarm Intelligence
- The role of Nature as source of examples and ideas to design new algorithms and multi-agent systems
- From observations to models and to algorithms
- Self-organized collective behaviors
- The role of space and communication to obtain self-organization
- Social communication and stigmergic communication
- Main algorithmic frameworks based on the notion of Swarm Intelligence: Collective Intelligence, Particle Swarm Optimization, Ant Colony Optimization
- Computational complexity, NP-hardness and the need of (meta)heuristics
- Some popular metaheuristics for combinatorial optimization tasks

Swarm Intelligence: what's this?

- Swarm Intelligence indicates a recent computational and behavioral metaphor for solving distributed problems that originally took its inspiration from the biological examples provided by social insects (ants, termites, bees, wasps) and by swarming, flocking, herding behaviors in vertebrates.
- *Any attempt to design algorithms or distributed problem-solving devices inspired by the collective behavior of social insects and other animal societies*. [Bonabeau, Dorigo and Theraulaz, 1999]

. . . however, we don't really need to "stick" on examples from Nature, whose constraints and targets might differ profoundly from those of our environments of interest . . .

Where does it come from?

- Nest building in termite or honeybee societies
- Foraging in ant colonies
- Fish schooling
- Bird flocking
- . . .

Nature's examples of SI

Fish schooling (©CORO, CalTech)

Nature’s examples of SI (2)

Birds flocking in V-formation (©CORO, Caltech)

Nature’s examples of SI (3)

Termites’ nest (©Masson)

Nature’s examples of SI (4)

Bees’ comb (©S. Camazine)

Nature’s examples of SI (5)

Swarm of killer bees (©S. Camazine)

Bees’ nest (©S. Camazine)

Nature's examples of SI (6)

Ant chain (©S. Camazine)

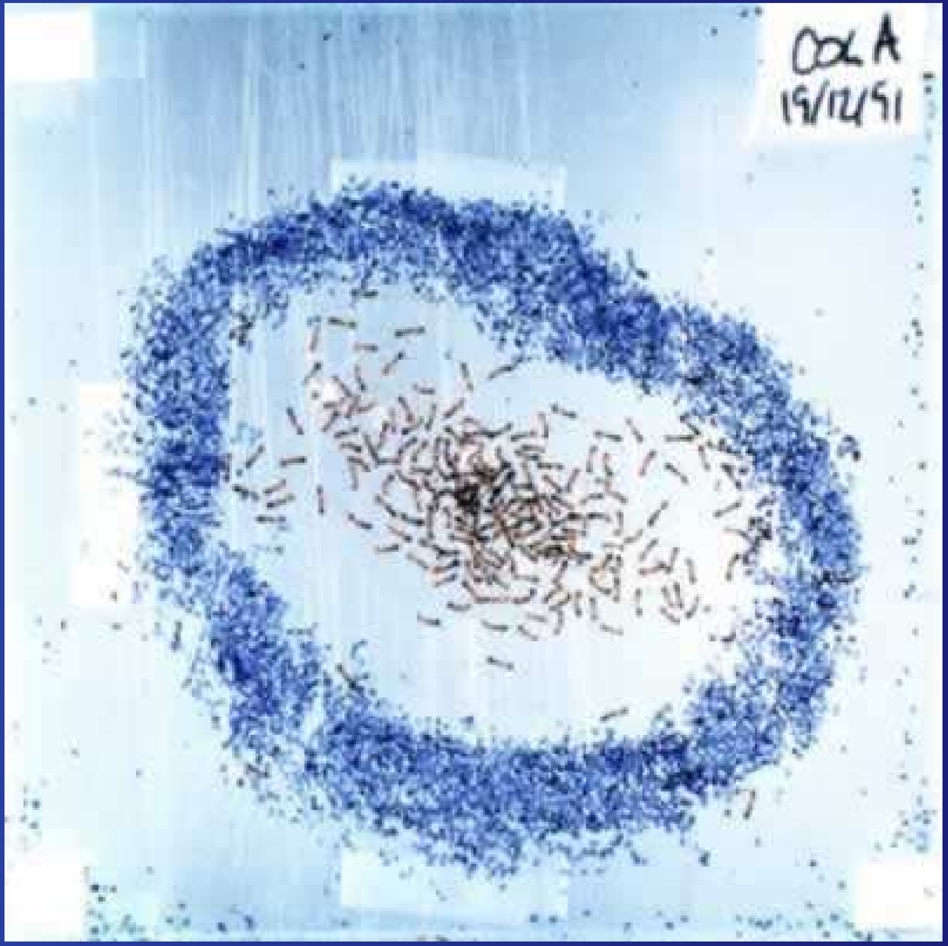

Ant wall (©S. Camazine)

Nature's examples of SI (7)

Wasps' nest (©G. Theraulaz)

Nature's examples of SI (8)

Bites of busy ant life:

- Leaf-cutting, breeding, chaining
- Food catering

What these behaviors have in common?

These are all intriguing, extremely fascinating behaviors that in spite of the specific diversity seem to be related to few invariant properties:

- Control is fully distributed among a number of individuals
- Communications among the individuals happen in a localized way
- System-level behaviors appear to transcend the behavioral repertoire of the single individual
- The overall response of the system is quite robust and adaptive with respect to changes in the environment

I had a dream . . .

. . . I can generate ***complexity out of simplicity***: I can put all the previous ingredients in a pot, boil them down and get good, robust, effective algorithms for my problems!

. . . it reminds me of alchemists . . .

It’s just about design choices

There’s no magic!

- Task complexity is a *conserved* variable
- ***Problem + Constraints + Optimization Criteria***: How do I solve it?
- Problems can be static, dynamic, online, offline, stationary, time-varying, centralized, distributed
- Algorithms can be monolithic, modular, distributed, parallel, adaptive. . .
- The final design choice is usually a rather obscure match between designer’s expertise, problem’s characteristics, constraints and targets

A Swarm Intelligence design is. . . ?

- Allocating computing resources to a number of relatively simple units (swarm?)
- No centralized control (not at all?)
- Units interact in a relatively simple and localized way
- . . . and I will get some useful global behavior

- What this has to do with Spatial Intelligence?

- The swarm *lives* distributed in some space
- *Communication* is a key aspect to get nonlinear behavior, and communication happens locally

BTW, what's *Intelligence*?

A more general definition of SI?

- Do we really need the explicit reference to insect/animal societies?
- Not really, however, let’s have first a closer look to some of these societies and to some of the interesting swarm behaviors that they can generate . . .

Few facts about Social Insects

- Social insects :
  - ***Ants***
  - ***Termites***
  - Some ***bees***
  - Some ***wasps***
- 10^{18} living insects (rough estimate)
- 2% of insect are social and most of them are eusocial
- 50% of all social insects are ants
- Total weight ants $\approx$ Total weight humans
  (one ant $1 \div 5$ mg)
- Ants are successfully around since 100 million years,
  Home sapiens sapiens only since 50,000 years

Ant colonies

- Ant colony size: from as few as 30 to millions of workers
- *Work division*:

Reproduction	$\longrightarrow$	Queen
Defense	$\longrightarrow$	Specialized workers
Defense	$\longrightarrow$	Soldiers
Food collection	$\longrightarrow$	Specialized workers
Brood care	$\longrightarrow$	Specialized workers
Nest brooming	$\longrightarrow$	Specialized workers
Nest building	$\longrightarrow$	Specialized workers

Some interesting collective behaviors

- Nest building and maintaining
- Division of labor and adaptive task allocation
- Discovery of shortest paths between nest and food
- Clustering and sorting (e.g., dead bodies, eggs)
- Structure formation (e.g., deal with obstacles)
- Recruitment for foraging (tandem, group, mass)
- Cooperative transport (e.g., food)

. . . and solitary ones: Ant navigation

- Depends on the sensorial capabilities of ant species as well as on the characteristics of the environment and function within the colony. Can make use of:

- Visual landmarks (use of memory and learning, encounters with colony mates)
- Chemical landmarks (pheromone)
- Compass-based (e.g., *Cataglyphis* desert ant uses light polarization)
- Dead-reckoning, path integration (calculation of the home vector)
- Correlated random walk

Let's go back to swarm behaviors

- The central question is: *How do social insects and other animals coordinate their actions in order to achieve amazing system-level behaviors?*
- Structures resulting from individuals' interactions develops by a process of ***Self-organization***
- BTW, amazing does not mean efficient. . .

Self-organization

- *Self-organization consists of set of dynamical mechanisms whereby structure appears at the global level as the result of interactions among lower-level components.*
  The rules specifying the interactions among the system's constituent units are executed on the basis of purely local information, without reference to the global pattern, which is an emergent property of the system rather than a property imposed upon the system by an external ordering influence [Bonabeau et al., 1997]

Characteristics of self-organization

Basic ingredients:

- Multiple interactions
- Amplifi cation of fluctuations and Randomness
- Positive feedback (e.g., recruitment and reinforcement)
- Negative feedback (e.g., limited number of available foragers)

Signatures:

- Creation of spatio-temporal structures (e.g., foraging trails, nest architectures, social organization)
- Multistability (e.g., ants exploit only one of two equivalent food sources)
- Existence of bifurcations when some parameters change (e.g., termites move from a non-coordinated to a coordinated phase only if their density is higher than a threshold value)

Is this definition satisfactory?

- There is no a unique and/or satisfactory definition of self-organization. The one provided here is not really a definition but rather a set of heuristic rules to design or spot self-organizing processes

- In more mathematical terms, I like the characterization given by Shalizi [Shalizi, 2001] who relates self-organization to the *statistical complexity* of the causal states of the process.

More on self-organization

- The statistical complexity $C_\mu(\mathcal{S})$ of the causal states $\mathcal{S}$ is defi ned as their *entropy measure* $H[\mathcal{S}]$ over the distribution μ of the inputs
- $C_\mu(\mathcal{S})$ is the average amount of bits that is retained about the input in the state set (inputs' partition)
- For a process the statistical complexity increases when self-organization is obtained: $C_\mu(\mathcal{S}_t) < C_\mu(\mathcal{S}_{t+T})$ (additional mathematical conditions are also considered)
- Intuitively, *when a number of units have reached organized coordination, it is necessary to retain more information about the inputs in order to make a statistically correct prediction*. In a coordinated phase precise information about all the participating units needs to be known in order to have a suffi cient picture of the system. That is, the entropy measure over the new causal states increases.

How is self-organization achieved?

- ***Communication is necessary***:
  - ***Point-to-point***: antennation, trophallaxis (food or liquid exchange), mandibular contact, direct visual contact, chemical contact, . . . unicast radio contact!
  - ***Broadcast-like***: the signal propagates to some limited extent throughout the environment and/or is made available for a rather short time (e.g., use of lateral line in fishes to detect water waves, generic visual detection, actual radio broadcast
  - ***Indirect***: two individuals interact indirectly when one of them modifies the environment and the other responds to the new environment at a later time. This is called ***stigmergy*** (e.g., pheromone laying/following, post-it, web)

Ant algorithms, Particle swarms and . . .

- Stigmergy has led to ***Ant Algorithms*** and in particular to ***Ant Colony Optimization (ACO)***
- Broadcast-like communication is related to schooling and flocking behaviors, that have inspired ***Particle Swarm Optimization***. In turn, neighbor broadcast is at the basis of ***Cellular Automata***, one of the early examples of swarm computation
- The use of all the three forms of communication encompasses more general systems showing collective organized behaviors (***COIN, immune system, cultural algorithms, neural system, human organizations, mobile ad hoc networks***,. . .)

A (tentative) more general definition of SI

- A swarm can be seen as a set of $N \gg 1$ communicating and distributed autonomous agents (e.g., ants, communication devices, computer processes) each engaged in one or more tasks, and with no or little centralized control
- If from the interactions among the constituents of the swarm results a process of self-organization that gives rise to interesting/useful behaviors at the system level, we can say that we are observing a phenomenon of ***Swarm Intelligence***
- Does ***Collective Intelligence*** sound better?
- Do we need restriction on aspects like:
  Nature-inspiration, short-range locality, agent simplicity, awareness of global task, homogeneity. . . ?
  $\longrightarrow$ Parameters

Some “new” examples from biology

- ***Immune system***: high diversity, mobility, distributed, dynamic, pipelined strategies, several communication strategies, multi-objective, learning, memory . . .

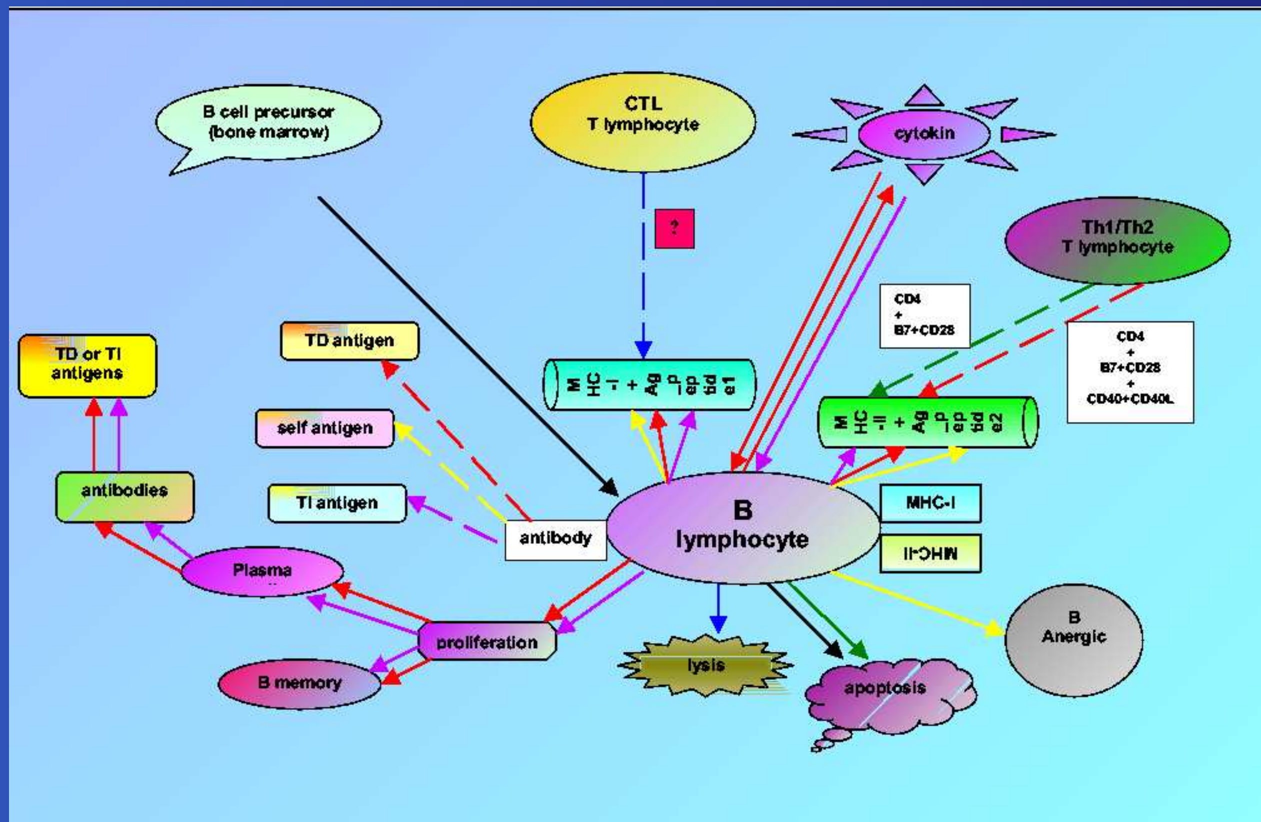

- ***Brains, slime molds, gene regulatory networks*** . . .

. . . and from “us”

- ***Routing in the Internet***: a system of distributed and adaptive controllers search online for good communication paths between computers

- ***Routing in mobile ad hoc networks***: each node is host and router, no infrastructures or centralized control, nodes might move and join and leave the network at any time, one shared communication medium, short range and noisy transmissions. Very dynamic and spatial-aware problem (more on Friday)

- ***Crowd control***: rush hours in Tokyo’s Shinjuku station, movies like ANTZ or Titanic

- In human organizations or teams, participants “usually” are aware of the global objectives. What’s the impact of this fact?

Collective robotics

- ***Collective robotics*** is attracting a lot of interest: groups of robots play soccer (RoboCup), unload cargo, cluster objects, self-assemble (Swarm-bots), and will “soon” participate to war :-(. . .

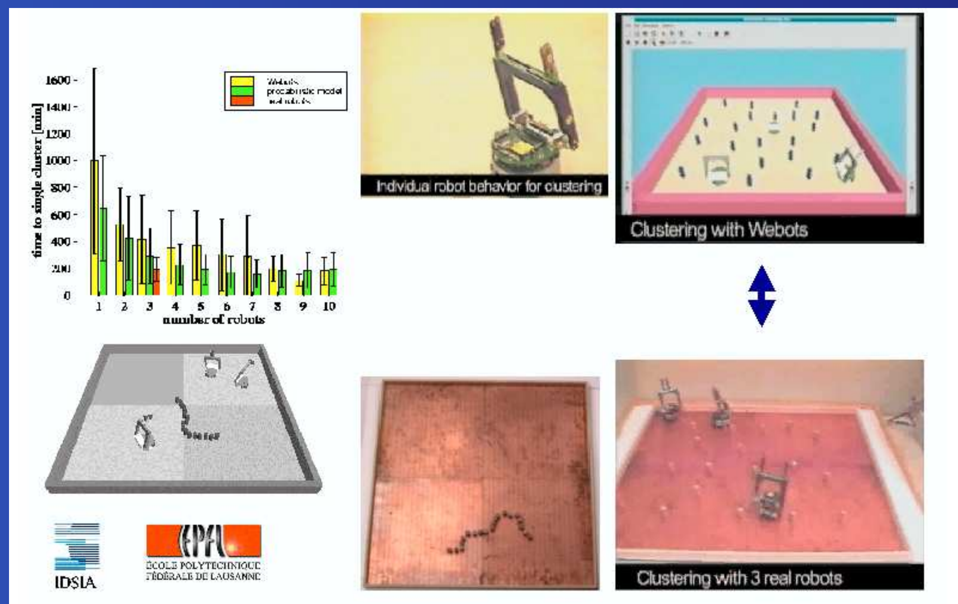

- Look at RoboCup (www.robocup.org) and Swarm-bots (www.swarm-bots.org)!

Which are the agents?

- Related question: Which is the right level of ***resolution*** to use when building a model of a (biological) system or designing an algorithm inspired by (biological) observations?

- When the model is about functional relationship and not concrete objects/beings the issue is even more critical (e.g., ant colony [Fewell, 2003])

- This is the main problem also when dealing with ***multi-agent*** (mechanistic) ***simulations*** (but phenomenological studies do not have better life, however). What's important and what can be leaved out in the simulator? Unfortunately most of the sensitivity studies reveal that "details" matters. However, science is precisely the continual search for and refinement of "good" models. . .

Back to the algorithmic frameworks. . .

- ***Ant Colony Optimization*** (ACO) and ***Particle Swarm Optimization*** (PSO) are the most popular instances of frameworks based on the original notion of SI (CA?)

- At the core of the design of ACO and PSO there is the specifi c way the agents communicate in the ***spatial environment***. These two ***optimization frameworks*** focus on two different ways of distributing, accessing and using ***information in the environment***

- In ACO and PSO agents are rather simple, since they ***do not learn at individual level***

- The agents in the ***Collective Intelligence*** (COIN) framework are reinforcements learners, therefore they can be arbitrarily complex. COIN's design focuses on ***generic multi-agent reinforcement learning***. Focus on the role of distributing and managing ***utility functions/values*** among the agents

- On a complexity scale, from the simplest: PSO, ACO, COIN

Othe related frameworks/keywords

Just names and buzzwords here:

Distributed Artificial Intelligence, Computational Economics, Multi-player Cooperative Game Theory, Evolutionary Computation (Population-based), Artificial Life, Statistical Physics, Markov Fields, Network Theory, Neural Networks, Traffic Theory . . .

(see [Wolpert and Tumer, 2000] for a list of references and comments)

COIN

The ***Collective Intelligence*** framework [Wolpert and Tumer, 2000] consists of:

- A large multi-agent system,
- where there is little to no centralized, personalized communication and/or control,
- there is a provided world utility function that rates the possible histories of the full system,
- each agent “runs” a reinforcement learning algorithm (microlearning).

COIN

- Central issues in COIN: ***how to map the world utility function into private utility functions for each of the agents?*** How the private utility functions can be designed so that each agent can realistically hope to optimize its function, and at the same time the collective behavior of the agents can optimize the world utility?

- The assignment of the rewards to the agents is a critical aspect in ***multi-agent reinforcement learning systems***

COIN

- COIN focuses on an ***inverse problem***: how to confi gure local dynamical laws and the management of the system-level utility in order to induce the desired global behavior.
  This is the ultimate dream of every engineer dealing with complex systems

- Fixing the agent characteristics and studying the response to different world utilities and distribution of them is also extremely useful: ***Economics***!

- COIN is a general mathematical framework. Not straightforward to understand. It points out where the problems are and provides formal tool to reason. However, does not provides straight or automatic design answers. It has been applied to routing and game problems. Not really popular, but worth to give a look at