

# Achieving Software Robustness via Large-Scale Multiagent Systems

Michael N. Huhns, Vance T. Holderfield, and Rosa Laura Zavala Gutierrez

University of South Carolina, Department of Computer Science and Engineering,  
Columbia, SC 29208, USA

{Huhns, Holderfield, Zavalagu}@engr.sc.edu

<http://www.cse.sc.edu/~huhns>

**Abstract.** This paper describes how multiagent systems can be used to achieve robust software, one of the major goals of software engineering. The paper first positions itself within the software engineering domain. It then develops the hypothesis that robust software can be achieved through redundancy, where the redundancy is achieved by agents that have different algorithms but similar responsibilities. The agents are produced by wrapping conventional algorithms with a minimal set of agent capabilities, which we specify. We describe our initial experiments in verifying our hypothesis and present results that show an improvement in robustness due to redundancy. We conclude by speculating on the implications of multiagent-based redundancy for software development.

## 1 Introduction

Computer systems are now entrusted with control of global telecommunications, electric power distribution, water supplies, airline traffic, weapon systems, and the manufacturing and distribution of goods. Such tasks are typically complex, involve massive amounts of data, affect numerous connected devices, and are subject to the uncertainties of open environments like the Internet. Our society has come to expect uninterrupted service from these systems. Unfortunately, when problems arise, humans are unable to cope with the complexity of the systems and the speed with which they must be repaired. Increasingly, the result is that critical missions are in jeopardy.

To cope with this situation, companies and researchers are investigating self-monitoring and self-healing systems, which detect problems autonomously and continue operating by fixing or bypassing the malfunction [25]. The techniques employed include redundant hardware, error-correction codes, and, most importantly, models of how a system *should* behave, so that the system can recognize when it *mis*behaves.

Such techniques can work well for hardware, but not for software, where having identical copies of a module provides no benefit. Software reliability is thus an unresolved problem. The amount of money lost due just to software errors is conservatively estimated at US\$40B annually. One alternative is to produce higher quality software. To achieve this, researchers are trying to define more

principled methodologies for software engineering [4]. They are also looking to new technologies, such as multiagent systems, and this leads to a natural interest in combining the latest software engineering methodologies with multiagent systems [22].

There are three ways that software engineering intersects multiagent systems:

- Multiagent systems can be used to aid traditional software engineering; e.g., in the form of agent-based or agent-supplemented CASE tools
- Traditional or new software engineering techniques can be used to build multiagent systems; e.g., the Unified Modeling Language (UML) has proven to be useful for conventional software, so Agent-Based UML (AUML) and similar efforts are underway to extend UML to support the development of agents
- Conventional software can be constructed out of agents, and software engineering techniques can be developed to aid in this endeavor.

The focus of this paper is on the last.

## 2 Background

### 2.1 Methodologies for MAS Modeling

Software engineering principles applied to multiagent systems have yielded few new modeling techniques, despite many notable efforts. A comprehensive review of agent-oriented methodologies is contained in Iglesias, et al. [12]. Many, such as Agent UML [23] and MAS-CommonKADS [11], are extensions of previous software engineering design processes. Others, such as Gaia [31], were developed specifically for agent modeling. These three have been investigated and applied the most. Other methodologies include the AAIL methodology [18, 17], MaSE [5], Tropos [26], Prometheus [24], and ROADMAP [15]. Because agents are useful in such a broad range of applications, software engineering methodologies for multiagent systems could be and should be a combination of efforts. This combination of principles and techniques will generally give a more flexible approach to fit a design team's particular expectations and requirements.

Agent UML, or AUML, is an extension of the Unified Modeling Language. UML was developed to unify object-oriented design techniques being randomly developed and utilized during the late eighties and early nineties. UML emphasizes things, relationships between things, and diagrams for grouping things and their relationships [2]. In 2000, Odell, et al. introduced an extension to UML, Agent UML, which includes agent interaction protocols. CommonKADS was developed for the design of knowledge-based systems (KBS); expert knowledge and a representation for problem domains form its basis [27]. Iglesias, et al. in 1997 proposed an extension for CommonKADS to include notations from several object-oriented modeling techniques (pre-UML). They also introduced seven knowledge models to take advantage of agent-oriented design. Gaia was formulated from an organization theory perspective. Its methodology is based on a system model consisting of roles, permissions, responsibilities, protocols, activities, liveness properties, and safety properties.

## 2.2 Benefits of an Agent-Oriented Approach

Multiagent systems can form the fundamental building blocks for software systems, even if the software systems do not themselves require any agent-like behaviors [14]. When a conventional software system is constructed with agents as its modules, it can exhibit the following characteristics:

- Agent-based modules, because they are active, more closely represent real-world things
- Modules can hold beliefs about the world, especially about themselves and others
- Modules can negotiate with each other, enter into social commitments to collaborate, and can change their mind about their results
- Modules can volunteer to be part of a software system.

The benefits of building software out of agents are [3, 10]

1. Agents enable dynamic composibility, where the components of a system can be unknown until runtime
2. Agents allow interaction abstractions, where interactions can be unknown until runtime
3. Because agents can be added to a system one-at-a-time, software can continue to be customized over its lifetime, even potentially by end-users
4. Because agents can represent multiple viewpoints and can use different decision procedures, they can produce more robust systems. The essence of multiple viewpoints and multiple decision procedures is redundancy, which is the basis for error detection and correction.

## 2.3 Bugs, Errors, and Redundancy

Hardware robustness is typically characterized in terms of faults and failures; equivalently, software robustness is typically characterized in terms of bugs and errors. Faults and bugs are flaws in a system, whereas errors and failures are the consequences of encountering the flaws during the operation or execution of the system. The flaws may be either transient or omnipresent. The general approaches for dealing with flaws are the same for both hardware and software:

- Fault/bug prediction and estimation
- Fault/bug prevention
- Fault/bug discovery
- Fault/bug repair
- Fault/bug tolerance or exploitation

Fault and bug estimation uses statistical techniques to predict how many flaws might be in a system and how severe their effects might be. For example, when Windows XP was released by Microsoft, it was estimated that it still contained 60,000 bugs, based on the rate at which its bugs were being discovered.

Bug prevention is dependent on good software engineering techniques and processes. Good development and run-time tools can aid in bug discovery, whereas bug repair and tolerance depend on redundancy.

Indeed, redundancy is the basis for most forms of robustness. It can be provided by replication of hardware, software, and information, and by repetition of communication messages. For years, NASA has made its satellites more robust by duplicating critical subsystems. If a hardware subsystem fails, there is an identical replacement ready to begin operating. The space shuttle has quadruple redundancy, and won't leave the ground without all copies functioning. However, software redundancy has to be provided in a different way. Identical software subsystems will fail in identical ways, so extra copies do not provide any benefit.

Moreover, code cannot be added arbitrarily to a software system, just as steel cannot be added arbitrarily to a bridge. When we make a bridge stronger, we do it by adding beams that are not identical to ones already there, but that have equivalent functionality. This turns out to be the basis for robustness in software systems as well: there must be software components with equivalent functionality, so that if one fails to perform properly, another can provide what is needed. The challenge is to design the software system so that it can accommodate the additional components and take advantage of their redundant functionality.

We hypothesize that agents are a convenient level of granularity at which to add redundancy and that the software environment that takes advantage of them is akin to a society of such agents, where there can be multiple agents filling each societal role [8]. Agents by design know how to deal with other agents, so they can accommodate additional or alternative agents naturally. They are also designed to reconcile different viewpoints.

Fundamentally, the amount of redundancy required is well specified by information and coding theory. Assume each software module in a system can behave either correctly or incorrectly. Then two modules with the same intended functionality are sufficient to detect an error in one of them, and three modules are sufficient to correct the incorrect behavior (by voting, or choosing the best two-out-of-three). This is exactly how parity bits work in code words. Unlike parity bits, and unlike bricks and steel bridge beams, however, the software modules can't be identical, or else they would not be able to correct each other's errors.

If we want a system to provide  $n$  functionalities robustly, we must introduce  $m \times n$  agents, so that there will be  $m$  ways of producing each functionality. Each group of  $m$  agents must understand how to detect and correct inconsistencies in each other's behavior, without a fixed leader or centralized controller. If we consider an agent's behavior to be either correct or incorrect (binary), then, based on a notion of Hamming distance for error-correcting codes,  $4m$  agents can detect  $m - 1$  errors in their behavior and can correct  $(m - 1)/2$  errors.

Fundamentally, redundancy must be balanced with complexity, which is determined by the number and size of the components chosen for building a system. That is, adding more components increases redundancy, but might also increase the complexity of the system. This is just another form of the common software engineering problem of choosing the proper size of the modules used to imple-

ment a system. Smaller modules are simpler, but their interactions are more complicated because there are more modules.

An agent-based system can cope with a growing application domain by increasing the number of agents, each agent's capability, the computational resources available to each agent, or the infrastructure services needed by the agents to make them more productive. That is, either the agents or their interactions can be enhanced, but to maintain the same degree of redundancy  $n$ , they would have to be enhanced by a factor of  $n$ .

To underscore the importance being given to redundancy and robustness, several initiatives are underway around the world to investigate them. IBM has a major initiative to develop autonomic computing—"a systemic view of computing modeled after the self-regulating autonomic nervous system." Systems that can run themselves incorporate many biological characteristics, such as self-healing (redundancy), adaptability to changing environments (reconfigurability), identity (awareness of their own resources), and immunity (automatic defense against viruses). An autonomic computing system will adhere to self-healing, not by "cellular regrowth," but by making use of redundant elements to act as replenishment parts. By taking advantage of redundant services located around the world, a better range of services can be provided for customers in business transactions.

For example, IBM's Tivoli Risk Manager is capable of automatically monitoring a network's health, protecting a network against attack, and healing a network in the event of attack. Among its autonomic features is a monitoring function referred to as the "heartbeat" that tracks so-called "keepalive" messages from third party security products and gives administrators an early warning about failures in their security infrastructure. Network devices use keepalive messages to communicate information about their online status and health. The heartbeat function is a way of actively monitoring the security products that Risk Manager manages. If a connection is lost, the heartbeat monitor issues an alert to the Risk Manager to take action or notify a human operator.

Exemplifying extreme redundancy in hardware, HP Labs has built a massively parallel computer, the Teramac, with 220,000 known defects, but it still yields correct results. As long as there is sufficient communication bandwidth to find and use healthy resources, it can tolerate the defects. Allowing so many defects enables the computer to be built cheaply.

The National Science Foundation has launched the Infrastructure for Resilient Internet Systems (IRIS) project, which is a five-year initiative to produce a robust, decentralized, and secure Internet infrastructure. The infrastructure will be developed using distributed hash table (DHT) technology, which can prevent all the data in a network from becoming vulnerable if one server crashes. Rather than centralizing the data in a single server, each server contains a partial list of the data's storage location; the challenge lies in developing a lookup algorithm that can locate data using the fewest possible steps. IRIS grew out of rising worries of the Internet's susceptibility to failure and attacks from viruses, worms, and possibly cyberterrorists.

## 2.4 N-Version Programming

N-version programming, also called dissimilar software, is a technique for achieving robustness first considered in the 1970's. It consists of  $N$  disparate and separately developed implementations of the same functionality. Although it has been used to produce several robust systems, it has had limited applicability, because (1)  $N$  independent implementations have  $N$  times the cost, (2)  $N$  implementations based on the same flawed specification might still result in a flawed system, and (3) the resultant system might have  $N$  times the maintenance cost (e.g., each change to the specification will have to be made in all  $N$  implementations).

## 2.5 Transaction Checkpointing, Rollback, and Recovery

Database systems have exploited the idea of transactions for maintaining the consistency of their data. A transaction is an atomic unit of processing that moves a database from one consistent state to another. Consistent transactions are achievable for databases because the types of processing done are very regular and limited.

Applying this idea to general software execution requires that the state of a software system be saved periodically (a checkpoint) so that the system can return to that state if an error occurs. The system then returns to that state and processes other transactions or alternative software modules. This is depicted in Figure 1.

There are two ways of returning to a previous state: (1) reloading a saved image of the system before the recently failed computation, or (2) rolling back, i.e., reversing and undoing, each step of the failed computation. Both of the ways suffer from major difficulties:

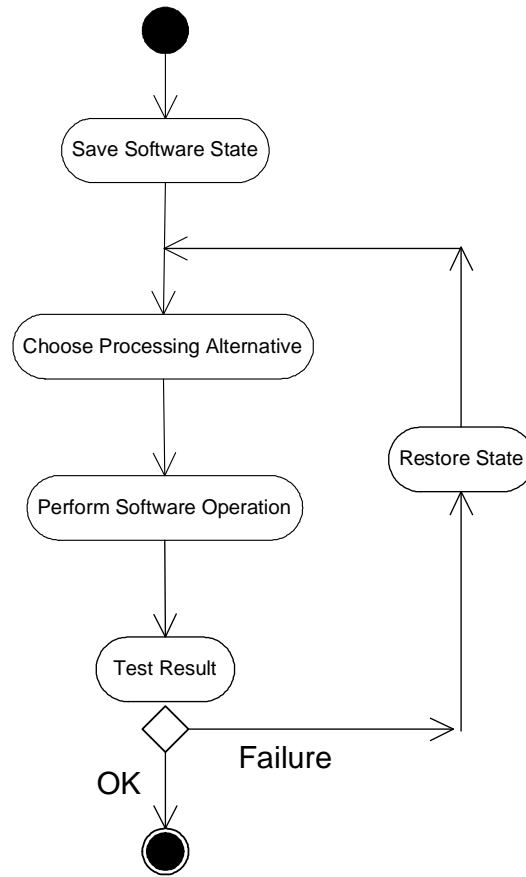
1. The state of a software system might be very large, necessitating the saving of very large images
2. Many operations cannot be undone, such as those that have side-effects. Examples of these are sending a message, which cannot be un-sent, and spending resources, which cannot be un-spent. Rollback is successful in database systems, because most database operations do not have side-effects.

## 2.6 Compensation

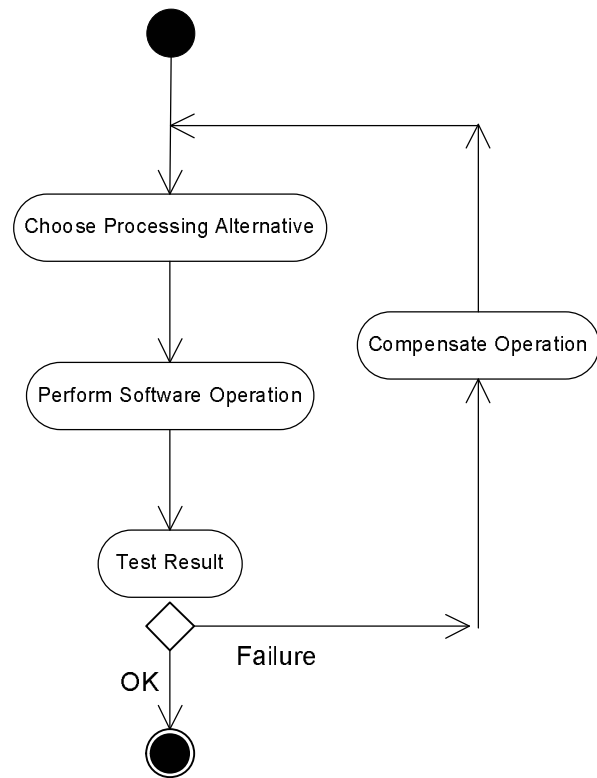
Because of this, compensation is often a better alternative for software systems. Figure 2 depicts the architecture of a robust software system that relies on compensation of failed operations.

## 3 Architecture and Process

Suppose there are a number of sorting algorithms available. Each might have strengths, weaknesses, and possibly errors. One might work only for integers, while another might be slower but be able to sort strings as well as integers.



**Fig. 1.** A transaction approach to correcting for the occurrence of errors in a software system



**Fig. 2.** An architecture for software robustness based on compensating operations



How can the algorithms be combined so that the strengths of each are exploited and the weaknesses or flaws of each are compensated or covered? In solving this in a general way, we hypothesize that the end result is an “agentizing” of each algorithm.

A centralized approach, as shown in Figure 3, would use an omniscient preprocessing algorithm to receive the data to be sorted and would then choose the best algorithm to perform the sorting. Each module’s characteristics have to be encoded into the central unit. The central unit can use a simplistic algorithm for determining best, based on known facts about each of the modules. The difficulties with this approach are (1) the preprocessing algorithm might be flawed and (2) it is difficult to maintain such a preprocessing algorithm as new algorithms are added and existing algorithms become unavailable. Also, only one module at-a-time executes, there is low CPU usage, and results are taken as-is when completed.

An improvement might be a postprocessing algorithm, as shown in Figure 4, that receives the results of all sorting algorithms and chooses the best result to be the output. Results have to be compared and voted on in order to determine the best. This approach is also centralized and suffers from a waste of CPU resources, because all algorithms work on the data. However, due to the comparison of outcomes, it is likely to produce better results.

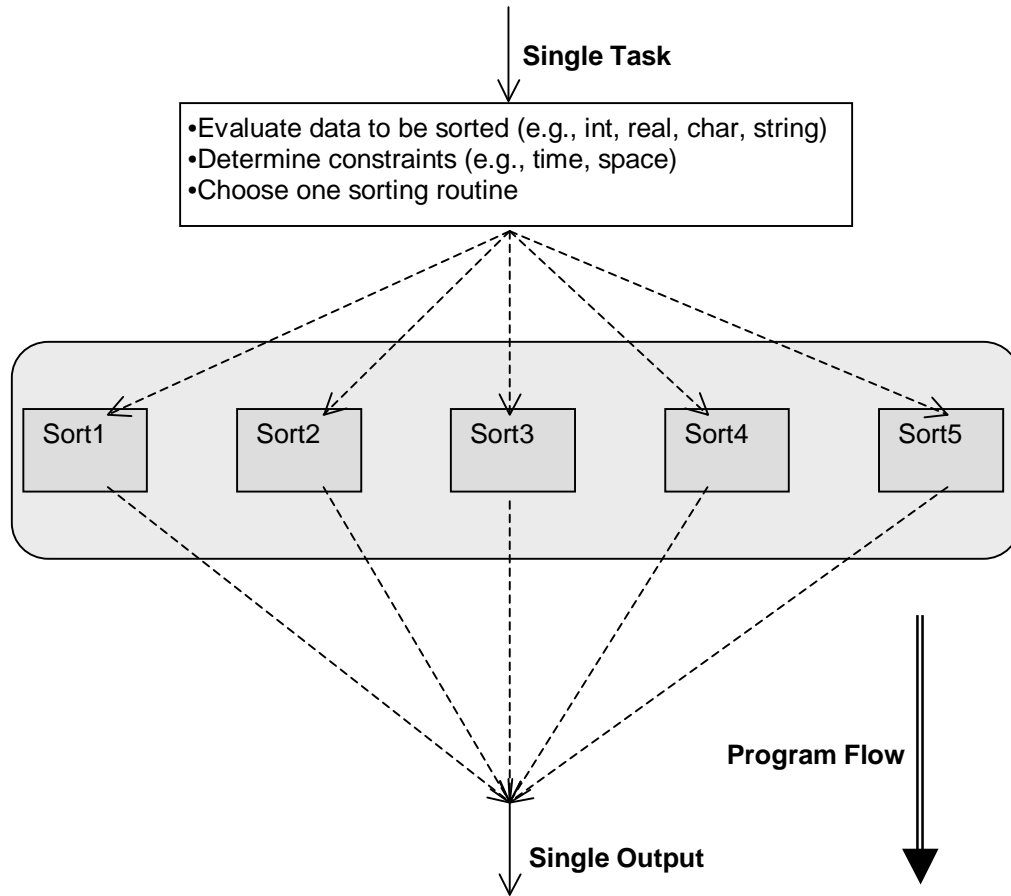
A combination of the pre- and postprocessing centralized systems could also be used. Since criteria are known about each module, a subgroup could be selected to perform the desired task based on known factors such as speed, time, and space. This subgroup would then have its results compared to determine the best results as above.

Certain modules will be selected every time the same set of circumstances come up, so a better way of developing a conventional system would then be to hard-wire these sets of circumstances to eliminate the need for a central intelligent filtering unit.

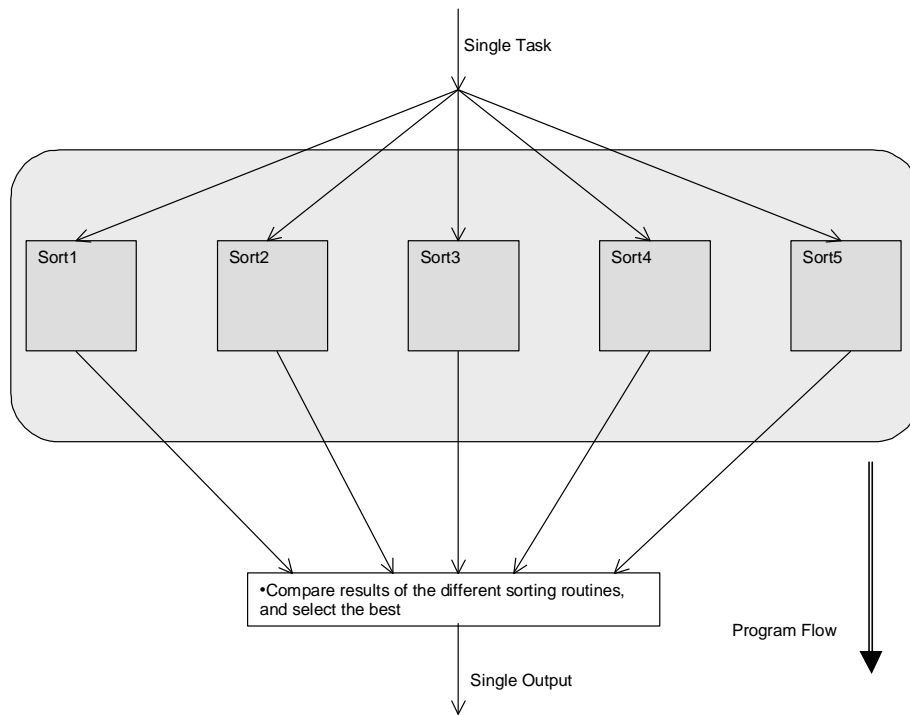
A fourth approach is a distributed solution, where the algorithms jointly decide which one(s) should perform the sorting, and if there is more than one qualified algorithm, they jointly decide on the best result. Conventional algorithms do not typically have such a distributed decision-making ability, so in this paper we investigate whether there is a generic capability that can be added to an algorithm to enable it to participate in a distributed decision. We also show that the result has the characteristics of a software agent.

### 3.1 Groups of Agents

Moving the administrative responsibilities from the central intelligent unit and distributing those responsibilities into the different modules creates a multiagent system. An agent in this system would have to know about itself: what it needs, what it can accomplish, and how. It would not have to know about its own algorithm to accept a task, just whether or not it can accomplish the task. Specifically, each algorithm must know



**Fig. 3.** Centralized architecture for combining  $N$  versions of a sorting algorithm into a single, more robust system for sorting, where a preprocessing algorithm chooses which sorting algorithm will execute



**Fig. 4.** Centralized architecture for combining  $N$  versions of a sorting algorithm into a single, more robust system for sorting, where a postprocessing algorithm chooses one result to be the output

- Something about its own algorithm, such as its time and space complexity, its input data structures, and its output data structures
- Something about other agents, such as their time and space complexity and reliability
- How to negotiate
- How to communicate
- How to compare results
- How to manage reputations and trust.

A transformation from the three conventional systems to an agent system would entail moving the responsibilities from a central unit, either a preprocessor or a postprocessor, to the individual agents.

An agent system transformed from a conventional system based on a preprocessor approach would have the agents themselves deciding on who would do the task placed before them. Since an agent knows about itself, it could bid for the task. This would create an auction environment to determine task assignment. This is an acceptable approach in a system with many competing agents, but the value of robustness based on reinforced redundant involvement is not instantiated in an auction environment.

An auction environment entails interaction among agents without any direct affect to the agents' abilities. Interaction without affecting another individual's problem solving capability can be best stated by Lorge and Solomon's "Model A" statement:

$$P = 1 - (1 - p)^r \tag{1}$$

The probability,  $P$ , of a group of individuals (agents) solving a problem is the probability that the the group size  $r$  contained at least one individual solver, given the probability,  $p$ , of a correct solution by an individual [21].

The idea here is that a group of agents can outperform an individual on a consistent basis [7]. Therefore the group, not competing individuals, should perform a task. The problem statement here is that a single task or set of task treated as a single unit is to be done by agents working redundantly in a group. A bidding system with a given set of mitigating circumstances will select the same agent under the same circumstances similar to the conventional system mentioned above.

An agent system transformed from a conventional system based on a centralized postprocessor would entail that each of the agents attempt the task and some type of voting mechanism (either with a voting factor or not) would be set up among the agents. A vote could be based on reputation only (more later) or on a majority rule vote based on the comparison of results. The communication overhead would/could be tremendous.

A key concept of groups is that a group can sometimes be wrong. A functional group basis will be more productive on average than the most productive member. Shapley and Grofman give the following example of five weather forecasters predicting whether it will rain or not on a given day. The decision is "yes, it will rain" or "no, it will not rain." The forecasters are given weights in

proportion to  $\log(p_i/(1-p_i))$ , where  $p_i$  is the probability of forecaster  $i$  making a correct decision. In their example, the forecasters are given the following weights: 0.9, 0.9, 0.6, 0.6, and 0.6. Assigning weights and determining a group decision by weighted voting generates a group probability of 0.927. This is higher than that of any one individual and also of unweighted voting, which has a group probability of 0.877 [28].

The notion of group is determined by the amount of communication among the group members [6]. Infrequent communication between individuals does not characterize the functioning of a group. These are casual relationships and do not hold to the group conscience. A group has a separate identity from each of its members, although the members do not cease to be individuals while in a group, because group members still have a personal responsibility to themselves, their own reputation, and their own desires.

A group is thought of as more than two members. One is an individual. Two is not considered a group since unless there is a consensus between the two; no group decision is possible without other influences, such as reputation of one member versus the other. But this is not a fair method, for whether right or wrong, the member with the better reputation will always see their reputation increase just because their choice was chosen. Outside influences need to mediate between groups that contain only two members.

There is an extreme number a group should hold, although this is not an exact number and should be applied to the different circumstances. A group generally stops being a group at about 10 to 15 members, after that the classification can be noted as an assembly (where members do more waiting around than not) or a mob (where members are out of control) [20].

The comparison of results can be done in several ways. The results can all be compared and a majority of exact outcomes would determine the results that are selected, where each of the agents have an equal chance at having their results selected. (For example, assume that there are four different results: A, B, C, and D. Five agents have result A, while two agents each have results B, C, and D. Therefore, result A is passed on as the accepted result.) This method is fine when there is a clear-cut winner, as above, but in the following case this methodology becomes cloudy. (For the scenario in the above example, assume three agents have result A and three agents have result B, while two agents have result C and D, respectively. A and B have the same number of votes, so an additional procedure must be used to choose between them.) A mitigating factor could be attributed to each agent based on information that the agent knows about itself, such as whether it completed the task or not, its time and space requirements, and the total number of runtime errors it has produced in the past.

A weighted value of importance could be calculated based on a agent's past performance, therefore a agent's chances of being selected is also based on past experiences.

## 4 Initial Experiments

We collected a number of sorting algorithms, each of them written by different people and therefore having different characteristics (such as input data type, output data type, and time and space complexity). For our experiments, we converted each algorithm into a sorting agent.

Each sorting agent is composed of a sorting algorithm and a wrapper for that algorithm. The wrapper knows nothing about the inner workings of the algorithm with which it is associated. It has knowledge only about the characteristics of its algorithm, such as the data type(s) it can sort, the data type it produces, its time complexity, and its space complexity. The sorting algorithms were written in Java and the wrappers in JADE [13].

The system sends data to be sorted to all the sorting agents. Their responsibility (as a group) is the sorting of the data, and they should be able to do this better than any one of them alone. Upon receiving data to be sorted, each agent determines whether or not it can sort it successfully (based on the type of the data and its own knowledge of what types it can sort) and if the agent can, it broadcasts an INFORM message to every other agent specifying its intention, along with a measure of performance for its algorithm (based on time and space complexity).

The decision of which agent (i.e., algorithm) to choose (among those that are capable of sorting the input data) is made in a distributed manner: upon receiving the INFORM messages from other agents, each agent compares its own performance measure against those received in the messages. If the agent has the best performance measure, it will run its algorithm and send the results back to the system. If it does not have the best performance measure, it will do nothing. Figure 5 depicts the AUML diagram of the protocol used. The system can be considered as the Initiator agent. Also, once they receive the data to be sorted from the system, the agents will wait for INFORM messages for only a limited amount of time; this avoids waiting infinitely long for messages from agents that either have problems sending a message or are not able to sort the data.

As can be seen, the current implementation of the wrappers is a front-ended approach. Agents themselves decide on who will do the task placed before them. Even with the simplicity of the current implementation, results showing improvement in robustness due to redundancy were obtained.

Table 1 summarizes the algorithms collected and information about them. The Additional Restrictions column shows restrictions not used by the default Wrapper to determine whether the agent can sort the data. Nevertheless, it is possible to cope with these cases by providing an implementation of a WrapperRestrictions interface, which only has one method that takes as input the data to be sorted and returns a Boolean value indicating whether the sorting algorithm can sort the data. This provides a way of customizing the wrapper for algorithms that need additional considerations for the whole system to perform better.

Finally, Table 2 provides a summary of the initial tests performed with the system.

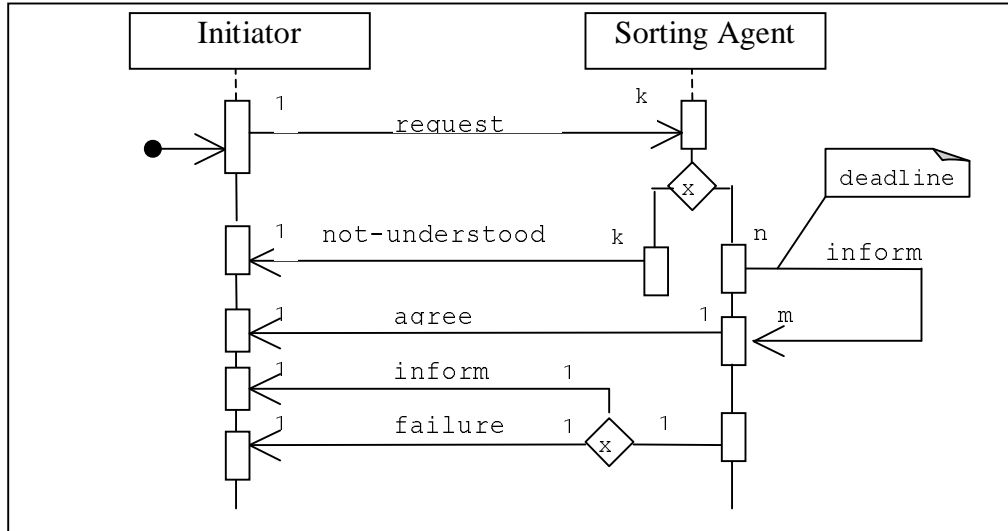


Fig. 5. AUML diagram of the agents' interactions

Table 1. Input data type and additional restrictions for the sorting algorithms available

Algorithm	Characteristics/Restrictions Used by Default Wrapper	Additional Restrictions
C.A.R Hoare's Quick Sort	Input data type: int array (positive and negative numbers accepted)	None
HeapSort	Input data type: int array (positive and negative numbers accepted)	None
QuickSort	Input data type: Byte array Short array Integer array Long array Float array Double array String array Char array	None
RadixSort	Input data type: int array	Only 10 inputs accepted

**Table 2.** Initial results for wrapping sorting algorithms with agents

Data Input	Algorithm Selected	Data Output	Comments
12, 45, 3, 2, 56	C.A.R Hoare's Quick Sort	2, 3, 12, 45, 56	
ann, john, sue, marie	QuickSort	ann, john, marie, sue	Only this algorithm could handle strings
9, 8, 7, 4, 3, 2, 12, 4, 5, 10	RadixSort	2, 3, 4, 4, 5, 7, 8, 9, 10, 12	Best performance but only worked for 10 inputs
3.54, 90, 23.4, 3.55, 60, 60.1	QuickSort	3.54, 3.55, 23.4, 60, 60.1, 90	Only this algorithm could handle reals

## 5 Group Decision Making

The system sends data to be sorted to the agent group (where all the sorting agents reside). For simplicity, we assume that no matter what kind of data is sent, there is at least one agent that can sort the data. Each agent receives the data and determines whether it can or cannot sort the data based on the data's type and its own knowledge of what kind of data it can sort. The agent will broadcast a 'yes' or 'no' on this subject. This first step will then determine the subgroup size. Because at least one agent can sort the data, the group size will range from 1 to  $N$ , where  $N$  is the total number of sorting agents in the group.

A group size of 1 is the trivial case. The agent who votes yes and 'hears' only negatives about sorting the input data will be the only one considered. This result is then passed on through the system and program flow proceeds as normal.

A group of size 2 to  $N$  is more complicated. When there is more than one positive concerning the ability to sort a data input, then a decision strategy must be put to use to determine which output to accept.

A voting strategy will be considered for this experiment. There are three basic voting incentives: the data, reputation, or mitigating factors. If the strategy involves the data, then direct comparisons need to be made about the data. An agent could compare another agent's result, which was positive, with its own. An array could be used to tabulate the results for each agent based on whether there is an exact match with itself or not.


For example, assuming only sortAgents 1, 2, 4, and 6 are capable of sorting the data, the array for sortAgent1 could be as follows:

$$\text{compare} = \text{Array}(\text{sortAgent2} \Rightarrow \text{yes}, \text{sortAgent4} \Rightarrow \text{no}, \text{sortAgent6} \Rightarrow \text{yes}) \quad (2)$$

Counting itself, sortAgent1 would receive three positives and one negative.

Comparisons could also be normalized according to the amount of similarity between different agent's results. Comparison could have a direct relationship between exact locations, a relationship between substrings, and/or a relationship between the number of location matches or exact substrings as in Figure 6.





**Fig. 6.** Partial comparison of different results could be based on exact location mappings (such as locations 1, 2, 3, 5, 6, 7, 9, 10), substring matching (such as 1-2-3, 5-6-7, and 9-10), and/or number of location matches or number of substring matches (eight and three, respectively).

CPU usage could be reduced if the result (yes or no) of a comparison is calculated only once, instead of by both agents. Logic could also aid in the reduction of CPU usage. If an agent A agrees with an agent B exactly, and agent A also agrees with agent C exactly, then agent B and agent C also agree. The sortAgent with the highest percentage of positives is selected. A tie can have two possibilities: (1) when the agents in the tie are a perfect match, then it does not matter which is chosen, or (2) when the agents are not a perfect match, then deciding is a problem.

If the strategy involves reputation, then a single sortAgent's vote is just its reputation factor; the result would be determined based on this factor alone. A tie here is a problem. Note that a reputation can be computed by many different means: e.g., neural nets, Bayesian networks, expert systems, or a simple normality based on successes. Reputation requires feedback for the system, either internally among the agents or externally from a user.

Mitigating factors, such as what the agent knows about itself, can affect a sortAgent's vote. Completion of the task has to be taken into consideration. Time and space constraints and runtime errors encountered during execution could be sorted in relative order of importance. A high number of runtime errors could affect a vote; a long delay or large space requirements could also affect a vote. A formula based on its time and space constraints and any runtime errors could be formed and used. A tie could be broken by hierarchically organizing each mitigating factor as above.

Combinations of these three basic voting incentives could be used. First, a formula based on mitigating factors would be used (since the presence of runtime errors is a very important fact); second, the data comparison would be made; and finally, a normalized reputation factor based on past successes in a group would be used. For example:

```

decisionVote = a * Completed? +
               b * number of runtime errors +
               c * time used +
               d * space used +
               e * exact comparisons +
               f * reputation (if there is one)

```

Variables a, b, c, d are mitigating factors, variable e is a data factor, and variable f is a reputation variable. The "best" algorithm is determined by

$$Performance \times Flexibility \times Reliability \quad (3)$$

where *Performance* is a function of time and space complexity, *Flexibility* is a function of how broad a range of input and output data structures the algorithm can handle, and *Reliability* is a measure of how well the algorithm can avoid runtime errors and exceptions.

## 6 Conclusion: Challenges and Implications for Developers

Producing robust software has never been easy, and the approach recommended here would have major effects on the way that developers construct software systems:

- It is difficult enough to write one algorithm to solve a problem, let alone  $n$  algorithms. However, algorithms, in the form of agents, are easier to reuse than when coded conventionally and easier to add to an existing system, because agents are designed to interact with an arbitrary number of other agents.
- Agent organizational specifications need to be developed to take full advantage of redundancy.
- Agents will need to understand how to detect and correct inconsistencies in each other’s behavior, without a fixed leader or centralized controller.
- There are problems when the agents either represent or use nonrenewable resources, such as CPU cycles, power, and bandwidth, because they will use it  $n$  times as fast.
- Although error-free code will always be important, developers will spend more time on algorithm development and less on debugging, because different algorithms will likely have errors in different places and can cover for each other.
- In some organizations, software development is competitive in that several people might write an algorithm to yield a given functionality, and the “best” algorithm will be selected. Under the approach suggested here, all algorithms would be selected.

Ultimately, the production of robust software will require that we understand the relationship between

- the social world as represented by humans and their physical environment, and
- the social world as represented by agents and other automated systems.

## References

1. Beugnard, Antoine, Jean-Marc Jezequel, Noel Plouzeau, and Damien Watkins: “Making Components Contract Aware,” *IEEE Computer*, Vol. 32, No. 7, July 1999, pp. 38–45.
2. Booch, Grady, James Rumbaugh, and Ivar Jacobson: *The Unified Modeling Language User Guide*. Addison-Wesley, Reading, MA, 1999.
3. Coelho, Helder, Luis Antunes, and Luis Moniz: “On Agent Design Rationale.” In *Proceedings of the XI Simposio Brasileiro de Inteligencia Artificial (SBIA)*, Fortaleza (Brasil), October 17–21, 1994, pp. 43–58.
4. Cox, Brad J.: Planning the Software Industrial Revolution. *IEEE Software*, (Nov. 1990) 25–33.

5. DeLoach, S.: "Analysis and Design using MaSe and agentTool." In *Proceedings of the 12th Midwest Artificial Intelligence and Cognitive Science Conference (MAICS 2001)*, 2001.
6. Dignum, Frank, Barbara Dunin-Keplicz, and Rineke Verbrugge: "Dialogue in team formation: a formal approach" In van der Hoek, W., Meyer, J. J., and Wittenveen, C., Editors, *ESSLLI99 Workshop: Foundations and applications of collective agent based systems*, (1999).
7. Hasling, John: *Group Discussion and Decision Making*, Thomas Y. Crowell Company, Inc. (1975).
8. Holderfield, Vance T. and Michael N. Huhns: "A Foundational Analysis of Software Robustness Using Redundant Agent Collaboration." In *Proceedings International Workshop on Agent Technology and Software Engineering*, Erfurt, Germany, October 2002.
9. Huhns, Michael N. and Vance T. Holderfield: "Robust Software," *IEEE Internet Computing*, vol. 6, no. 2, March-April 2002, pp. 80–82.
10. Huhns, Michael N.: "Interaction-Oriented Programming." In *Agent-Oriented Software Engineering*, Paulo Ciancarini and Michael Wooldridge, editors, Springer Verlag, Lecture Notes in AI, Volume 1957, Berlin, pp. 29-44 (2001).
11. Iglesias, C. A., M. Garijo, J. C. Gonzales, and R. Velasco: "Analysis and Design of Multi-Agent Systems using MAS-CommonKADS." In *Proceedings of the AAAI'97 Workshop on agent Theories, Architectures and Languages*, Providence, USA, 1997.
12. C. Iglesias, M. Garijo, and J. Gonzalez: "A survey of agent-oriented methodologies." In J. Muller, M. P. Singh, and A. S. Rao, editors, *Proceedings of the 5th International Workshop on Intelligent Agents V: Agent Theories, Architectures, and Languages (ATAL-98)*. Springer-Verlag: Heidelberg, Germany, 1999.
13. *JADE: Java Agent Development Environment*, <http://sharon.csel.it/projects/jade>.
14. Jennings, Nick R.: "On Agent-Based Software Engineering" *Artificial Intelligence*, 117 (2) 277–296 (2000).
15. Juan, T., A. Pearce, and L. Sterling: "Extending the Gaia Methodology for Complex Open Systems." In *Proceedings of the 2002 Autonomous Agents and Multi-Agent Systems*, Bologna, Italy, July 2002.
16. David Kalinsky: "Design Patterns for High Availability." *Embedded Systems Programming* (August 2002) 24–33.
17. Kinny, David and Michael Georgeff: "Modelling and Design of Multi-Agent Systems," in J.P. Muller, M.J. Wooldridge, and N.R. Jennings, eds., *Intelligent Agents III — Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages*, Springer-Verlag, Berlin, 1997, pp. 1–20.
18. Kinny, D., M. Georgeff, and A. Rao: "A Methodology and Modeling technique for systems of BDI agents." In *Proceedings of the 7th European workshop on modeling autonomous agents in a multi-agent world*, LNCS 1038, pp. 56–71, Springer-Verlag, Berlin Germany, 1996.
19. Laddaga, Robert: "Creating Robust Software through Self-Adaptation," *IEEE Intelligent Systems*, Vol. 14, No. 3, May/June 1999, pp. 26–29.
20. Light, Donald, Suzanne Keller, and Craig Calhoun: *Sociology* Alfred A. Knopf/New York (1989).
21. Lorge, I. and H. Solomon: "Two models of group behavior in the solution of Eureka-type problems." *Psychometrika* (1955).
22. Nwana, Hyacinth S. and Michael Wooldridge: "Software Agent Technologies." *BT Technology Journal*, 14(4):68-78 (1996).

23. Odell, J., H. Van Dyke Parunak, and Bernhard Bauer: "Extending UML for Agents." In *Proceedings of the Agent-Oriented Information Systems Workshop*, Gerd Wagner, Yves Lesperance, and Eric Yu eds., Austin, TX, 2000.
24. Padgham, L. and M. Winikoff: "Prometheus: A Methodology for Developing Intelligent Agents." In *Proceedings of the Third International Workshop on Agent-Oriented Software Engineering*, at AAMAS 2002. July, 2002, Bologna, Italy.
25. Paulson, Linda Dailey: "Computer System, Heal Thyself," *IEEE Computer*, (August 2002) 20-22.
26. Perini, A., P. Bresciani, F. Giunchiglia, P. Giorgini, and J. Mylopoulos: "A knowledge level software engineering methodology for agent oriented programming." In *Proceedings of Autonomous Agents*, Montreal CA, 2001.
27. Schreiber, A. T., B. J. Wielinga, and J. M. A. W. Van de Velde: "CommonKADS: A comprehensive methodology for KBS development," 1994.
28. Shapley, L. S. and B. Grofman: "Optimizing group judgmental accuracy in the presence of interdependence" *Public Choice*, 43: 329-343 (1984).
29. Swap, Walter C., et al.: *Group Decision Making*, SAGE Publications, Inc., Beverly Hills, London, New York (1984).
30. Tambe, Milind, David V. Pynadath, and Nicolas Chauvat: "Building Dynamic Agent Organizations in Cyberspace," *IEEE Internet Computing*, Vol. 4, No. 2, March/April 2000.
31. Wooldridge, M., N. R. Jennings, and D. Kinny: "The Gaia Methodology for Agent-Oriented Analysis and Design." *Journal of Autonomous Agents and Multi-Agent Systems*, 2000.