MannTall - A Rescue Operations Assistant

Martin Brooks, Steinar Carlsen, Atle Honne

Center for Industrial Research
Box 124 Blindern
0314 Oslo 3, Norway

Abstract

MannTall is a decision support system for deployment of rescue craft during an emergency on an offshore oil platform. MannTall uses heuristic counting and estimation techniques to compute best- and worst-case scenarios for the locations of the platform crew. MannTall is fully operational as a part of Saga Petroleum's safety and emergency preparedness procedures.

Use of MannTall in Large Scale Rescue Operations

MannTall is a stand-alone system for keeping track of personnel during an emergency on an offshore oil platform. MannTall is used at Saga Petroleum's Emergency Operations Room at Forus, on the west coast of Norway. MannTall provides a running analysis of the whereabouts of the platform crew in the form of upper and lower numerical bounds and the possible identities of crew members on the platform, in the sea, and on the various rescue craft. MannTall's input is a stream of messages regarding sightings and transfers, for example: "Abel , Berg, and one unidentified person have been picked up from the man-overboard boat by the helicopter". MannTall's analysis is used as decision support for deployment of rescue craft. The name "MannTall" is Norwegian for "census" or "roll call", literally as "person-count".

MannTall is meant for use in accidents serious enough to warrant evacuation of the platform, such as an explosion or blowout. Such situations involve more than a hundred people, rapid mobilization of rescue craft, and a high volume of radio and telephone communication. Due to the large amount and diversity of information involved, Saga's onshore crisis management team faces an extremely complex decision making task which they must perform under intense stress and time pressure. Their first priority is platform crew welfare.

Before deployment of MannTall in early 1988, personnel tracking was carried out by posting slips of paper on a bulletin board organized by peoples' names. Each slip contained information about a sighting or transfer of named crew members. Several people were responsible for maintaining and interpreting the information. MannTall is a direct substitution for the bulletin board and its manual analysis; MannTall logs the information and automatically provides the analysis. MannTall is operated by one or two people.

MannTall has three advantages: The first is MannTall's ability to provide excellent best- and worst-case estimates in the presence of ambiguous and incomplete information, including reports about unidentified crew members. Formerly, estimates were based solely on intuition and were limited by the complexity of the situation. The second advantage is that MannTall facilitates on-the-fly change of the persons responsible for personnel tracking. Since all message history and possible interpretations are on the computer screen, instead of in people's heads and their personal notes, the transfer of responsibility can occur efficiently and without loss of information. The third advantage is that fewer people are needed to carry out personnel tracking.

The value of MannTall's estimation ability is measured by the increased probability of saving a life. Test scenarios have demonstrated that MannTall can discover the possibility of an unrescued man in the sea in situations too complex for intuitive human analysis. MannTall is successful in Saga's high-stress emergency operations environment because its analytic power is made available via a simple interface.

Problem Definition

During the first hours of an offshore emergency, sightings and transfers of people are reported to Saga's Emergency Operations Room. The information comes from a variety of sources, including the government's regional Rescue Coordination Center, the police, and companies employing people on the platform, as well as direct monitoring of offshore radio communications. Problems arise in interpreting the information because:

- The same event may be reported several times.
- A message may be distorted when retransmitted through several channels.
- The contents of different messages may be conflicting.
- Some events may never get reported.
- Some messages refer to a single event, while others may refer to the results of several events.

It is extremely difficult for Saga's emergency operations staff to discover all possible interpretations of the messages, to maintain an overview of their combined effects, or to know which additional information would reduce the uncertainty. This is what MannTall does.

MannTall's Capabilities

MannTall's initial state is a list of all persons on board the platform. (This data is maintained by an independent system connected to MannTall.) As the rescue operation evolves, messages regarding the transfer and current location of people are entered. MannTall computes upper and lower numerical bounds and the possible identities of the people at each place, including the platform, sea, and the various rescue craft. Furthermore, MannTall generates questions to the user asking about the ambiguities inherent in the messages received so far. Answers to these questions will tighten the bounds. The user enters the answers as the information comes in, and in the meantime may use the questions as a what-if mechanism for exploring the possibilities.

MannTall's user interface has been constructed with efficiency and clarity as the main objectives - both are necessary in the high-stress operational environment in which the system is used. All input operations are menu and mouse driven. The available information can be presented using several different viewpoints; it is always up to date.

Figure 1 shows MannTall's main screen layout. The "Message Window" on the right shows a summary of the input messages in chronological order. Each message describes either a reported transfer of people from place to place or the reported status of the people at a place. For example; the message labeled T4 reports two unidentified people being picked up from the sea by the standby boat. The message labeled T7 reports one person being transferred from the standby boat to a helicopter; in this case

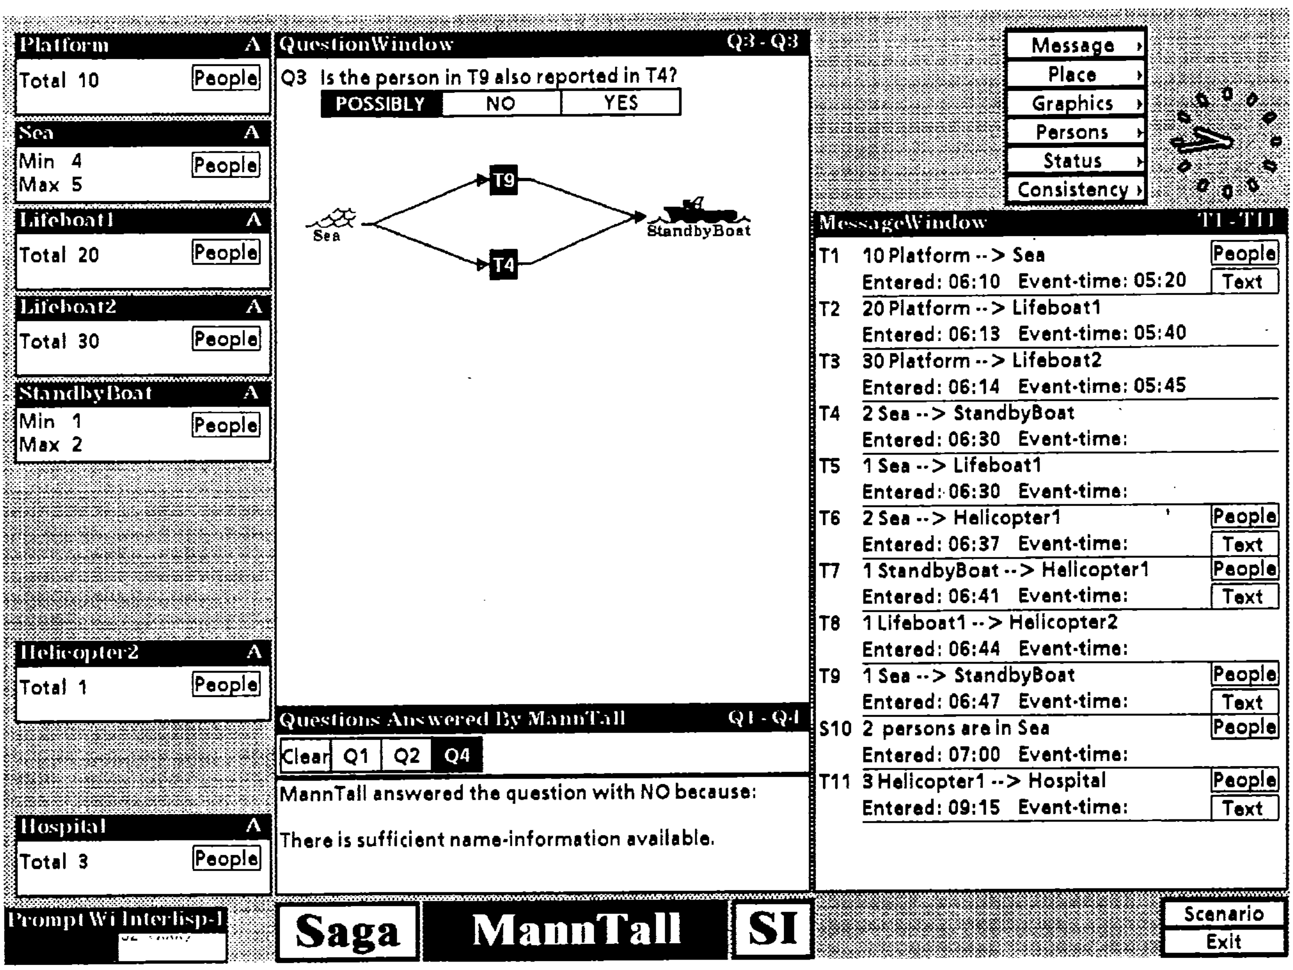

Figure 1: MannTall's main screen

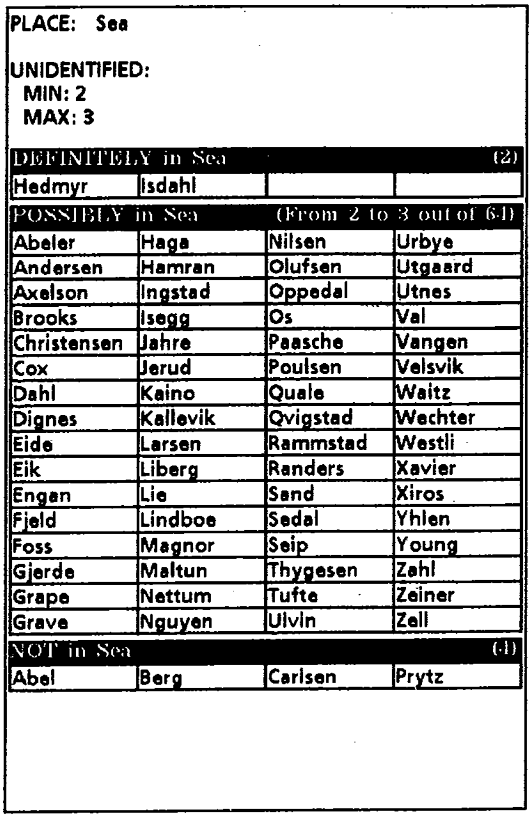

Figure 2: Window showing who can be in the sea.

the person's name is known and can be examined by clicking the mouse in the "people" button to the right of the message.

The small windows on the left of Figure 1 correspond to places where people can be, and give the maximum and minimum number at each place. The title bar of each window gives the place. For example, the "sea" window shows that there are 4 - 5 people in the sea. Clicking the mouse in the "people" button in this window pops up the list in Figure 2. This list shows the names of the two people definitely known to be in the sea in the topmost part, those possibly in the sea in the middle part, and those definitely known to not be in the sea in the bottom part.

The center "Question Window" of Figure 1 shows MannTall's questions to the user. The shown question asks whether the person reported in message T9 could be the same as one of the people reported in message T4. Note that T9 and T4 both report transfers from the sea to the standby boat, as shown in the graphic below the question. The answers to this question are "possibly", "no", and "yes" and are presented as three buttons under the question. The answer shown is "possibly"; this uncertainty accounts for the range of 4 - 5 people in the sea and 1 - 2 people in the standby boat. The answer can be changed at any time by clicking in the "yes" or "no" buttons; new calculations are immediately performed. For example, clicking "no" gives 4 people in the sea (i.e. min = max) and 2 people in the standby boat.

Below the Question Window is a window describing other questions that were automatically answered by MannTall. In this example the user has clicked in "Q4" and MannTall is explaining the basis for its automatically generated answer. The user may also pull up the text of question Q4 and change MannTall's answer, if necessary.

MannTall is a relatively mature product; the computational core is augmented by a large amount of supporting functionality, including serial-line connection to Saga's persons-on-board database, status reports, information logging facilities, consistency analysis, several forms of graphical presentation, and an extensive scripting facility for development and analysis of test scenarios.

Development and Fielding

MannTall was originally commissioned in order to demonstrate AI technology within Saga Petroleum. Safety was chosen because it is an open area with no sensitive knowledge or information, and in this area oil companies normally share their experience and results. Identification of the exact problem to be solved took several months.

MannTall was prototyped during 1986 by the authors in close cooperation with Saga. The initial work focussed on the core functionality and the user interface. As early as autumn 1986, experienced emergency preparedness managers at Saga recognized MannTall as a significant improvement over existing procedures.

In 1987 Saga began using MannTall as part of its regular simulated-emergency safety exercises, validating MannTall's usefulness and usability. In mid-1987 Saga ordered an operational version, which was delivered in late 1987. The system has been further refined based on user experience during 1988.

Use during a major exercise in 1988 had very satisfying results: throughout the entire exercise the overview of the platform crew's whereabouts was as complete as possible given the available information, a situation seldom achieved before MannTall.

MannTall is installed in Saga's Emergency Operations Room, and is part of Saga's formally defined procedures for reacting to an emergency. Fortunately, serious platform accidents are rare; MannTall has been used only once in a real emergency. On January 20, 1989, the Treasure Saga platform in the North Sea was exposed to a potential blowout and was partially evacuated. MannTall was used continuously for 48 hours while people were ferried back and forth between the platform and land by helicopter. Saga reports that they were very satisfied with MannTall's performance.

MannTall was developed and deployed on the Xerox 1186. The majority of MannTall's four man-year programming effort went to "productization" of the original prototype; three man-years of non-programming effort went into test, evaluation, and integration into Saga's procedures. Development costs were about $600K.

MannTall is generically designed so that it can be adapted to other emergency organizations.

How MannTall Works

MannTall is a based on a "reported-transfer graph" having one node for each place and one directed arc for each reported transfer of people from place to place. Each reported transfer is labeled with pertinent information, including the time of the report, the number of unidentified people, and the names of those identified. MannTall recognizes certain types of ambiguities in the reported-transfer graph; these ambiguities are the basis for multiple interpretations of the reported transfers.

MannTall's core is the computation of minimum and maximum numbers, and possible identities, of the people at each place in the reported-transfer graph. In theory one could find this information by creating one copy of the graph for each interpretation of the reported transfers, however combinatorics make this impractical. Instead, MannTall keeps the ambiguity "packaged" in the single-graph representation, and uses heuristic counting and estimation methods to interpret the reported transfers. This is a novel and surprisingly difficult problem; its solution is the the primary technical innovation in MannTall. The important property of these estimations is that they bracket the true value, in particular that they never rule out an actual possibility.

MannTall recognizes reference ambiguities in the reported-transfer graph. For example, suppose there have been two reports of three unidentified people falling into the sea from the platform. If these reports actually refer to the same people, then only three people are in the sea, whereas if they refer to completely different people, then a total of six are in the sea, and if one person was common to both, then a total of five are in the sea. A reference ambiguity may also involve a chain of transfers; for example, suppose two men are reported picked up from the sea by the standby boat, and then the standby boat reports the helicopter picking up one of the men, and thirdly the helicopter reports picking up a man who was in the sea. The third report may or may not refer to one of the men in the first two reports. Figure 3 shows the reported-transfer graph for these two examples.

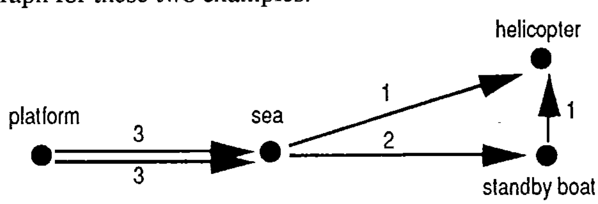

Figure 3: Two types of ambiguity in the reported transfer graph. There are 0-4 people in the sea, 1 person picked up from the sea in the standby boat, 1-2 people in the helicopter who have been in the sea.

MannTall uses a combination of techniques to compute the range of numbers and possible identities of the people at each place. The structure of the total solution is shown in Figure 4, where the arrows show the relations between subproblems. Some of the individual subproblems are briefly described in the following paragraphs; the point of these descriptions is only to give the reader a feeling for the types of solutions.

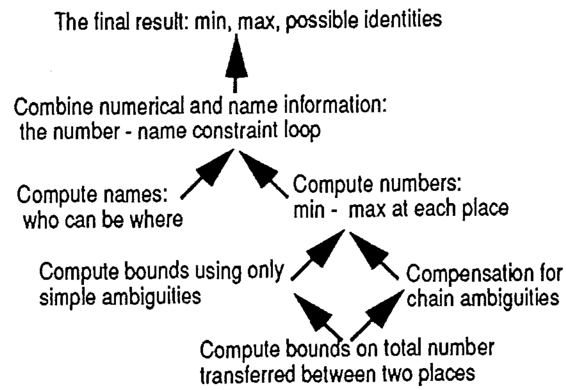

Figure 4: Structure of MannTall's heuristic counting and estimation techniques

Compute bounds on total number transferred between two places

Problem: Given a set of reported transfers, each from place A to place B, find the minimum and maximum total number of people transferred from A to B.

Note: The possibility of reported transfers referring to the same people is dependent on the user's answers to MannTall's questions.

Solution: A solution that ignores the names of identified people in the transfers is developed first. Tighter bounds are then derived by taking names into account. The solution has the form of a recursive counting procedure around a set of constraints.

Compute numbers: min - max at each place

Problem: Find the maximum and minimum number of people at a particular place.

Solution: The number of people at a place is figured as the difference between the number coming in and the number going out. Uncertainty ranges, [min, max], are added in order to get the totals in and out, and subtracted in order to get the difference between the total in and total out:

$$[min_1 , max_2] + [min_2, max_2] = [min_1 + min_2 , max_1 + max_2]$$
$$[min_{in} , max_{in}] - [min_{out}, max_{out}] = [min_{in} - max_{out} , max_{in} - min_{out}]$$

If there are no chains, then this is straightforward, as in Figure 5a. With chain ambiguities, an additional technique is needed to recognize the possibility of counting the same person twice, as in Figure 5b.

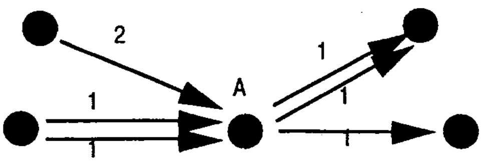

Figure 5a: Without chains, the number of people at A is difference between the totals coming in and going out from other places; i.e.

([2,2] + [1,2]) - ([1,1] + [1,2]) = [3,4] - [2,3] = [0,2]

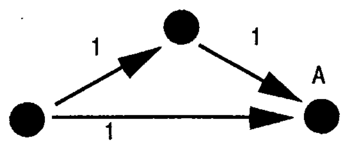

Figure 5b: With chains, one must be careful not to count the same person twice. The number of people at A is [1,2].

Compensation for chain ambiguities

Problem: Chains of transfers create ambiguities where two people reported arriving at (or leaving from) the same place from (to) different places may in fact be the same person. Given a network of chains ending at a place, compute the minimum and maximum number of people that could be counted more than once.

Solution: Overlapping and nested chains make the problem particularly difficult, because it is not sufficient to analyze them individually, as illustrated in Figure 6. The estimate is made by analyzing them separately while constructing a table coding the effects of overlap according to a topological classification of overlap types.

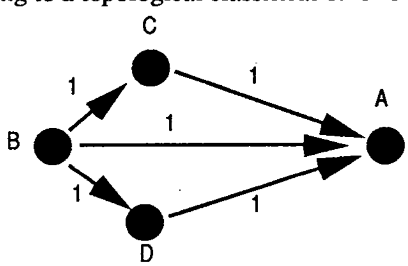

Figure 6: Overlapping chains. The reported transfer from B to A is ambiguous with both the B-C-A chain and the B-D-A chain. The chain ambiguities cannot be analyzed separately, since the person reported from B to A cannot simultaneously participate in both chains.

Combine numerical and name information: the number - name constraint loop

Problem: Numerical bounds on the number of people at a place are derived as described above. Lists of names of people definitely at, and possibly at the place are derived by other methods. These two estimates do not necessarily agree, and need to be combined.

Solution: The numerical estimates and name data constrain each other according to a set of rules which are iterated forward until the results are stable. Each iteration tightens the bounds.

Acknowledgements

The authors wish to thank Birger Nordberg, Tor Stein Ølberg, Frode Leraand, and Bo Brännström at Saga Petroleum, as well as Freddy Krogh and Morten Irgens at the Center for Industrial Research.

The aim of the Department for Knowledge Based Systems at the Center for Industrial Research is to make AI technology operational in real-world problems. Our efforts are directed towards problems of high complexity within the industrial sector. Our product is custom designed software based on unique generic modules and know-how within AI.

Martin Brooks received a B.S. in mathematics from MIT in 1974 and a Ph.D. in computer science from Stanford University in 1980. His current interests are AI, robotics, mathematical theory of computation, and neural nets.

Steinar Carlsen received an M.Sc. in computer science from the Norwegian Institute of Technology in 1985. His current interests are AI, knowledge engineering, CASE tools and methodology, and production planning. He led the MannTall development project in 1987-8.

Atle Honne received an M.Sc. in applied physics from the Norwegian Institute of Technology in 1973. His current interests are AI and image processing and analysis. He is the manager of the Department for Knowledge Based Systems at the Center for Industrial Research.