Distribution Models for Network Management Functions

Frank Strauß

Computer Science Department
Technical University Braunschweig
Bültenweg 74/75
38106 Braunschweig, Germany
strauss@ibr.cs.tu-bs.de

June 2000

Abstract

To ensure scalability, efficiency and flexibility of network and systems management platforms of today's growing networks, distributed management systems are strongly required. A quite commonly proposed simple taxonomy suggests a distinction between centralized, weakly distributed, strongly distributed and cooperative management distribution models. However, it is not yet clearly stated which distribution models are best suited for certain categories of management problems.

This article presents a number of criteria of given management problems. Multiple combinations of these criteria lead to some reasonable characteristics of management systems that are suitable to solve the management problem. In turn, typical combinations of these characteristics make up the four aforementioned management distribution models.

1 Distributed Management

The early network management approaches at the end of the 1980s were based on a strictly centralized architecture consisting of a central network management station and a number of agents supplying the management information.
However, the demand for enhanced management systems is growing regarding to their flexibility, robustness, efficiency and scalability: the size and bandwidth of managed networks is increasing constantly, along with the heterogeneity of systems, network technologies, and services. Meanwhile, the computational power of managed systems increases significantly. These facts are generally addressed by the concepts of distributed management. It is based on the delegation of management tasks to dual-role entities that are called *mid-level managers*, so that the network bandwidth can be reduced and the need for CPU power and storage requirements can be distributed among these mid-level managers. Furthermore, robustness can be increased through the use of redundancy, and flexibility can be improved by dynamic delegation of code and even processes [1].

Before administrative advantages can be made out of these circumstances, a given management problem has to be well analyzed. The following section presents a number of relevant management problem *criteria*. Some rules, presented in section 3, combine these criteria and lead a number of management system *characteristics* suitable for the given situation. In turn, the most likely combinations of these characteristics lead to four management *distribution models*, even though there is no strict discrimination: *centralized*, *weakly* and *strongly distributed*, and *cooperative management* [2, 3]. This correlation is described in section 4. The process of this analysis is shown in figure 1.

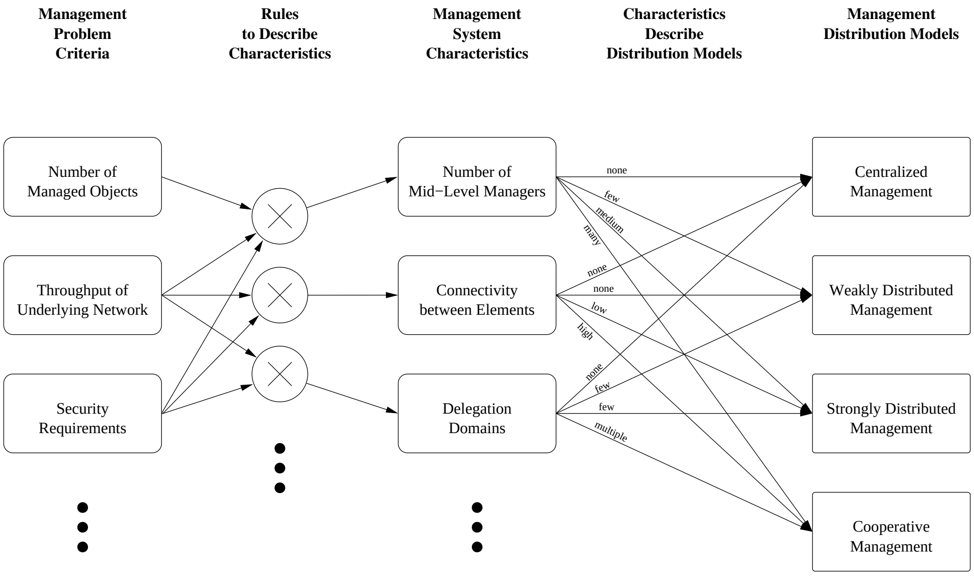

Figure 1: From a management problem to management system characteristics to management distribution models.

2 Criteria of Management Problems

This section lists eight criteria of management problems, that should be considered for a given situation to conclude with a set of significant management system characteristics. Some of these criteria are related and lead to common characteristics while others are completely independent [4].

Most of these criteria descriptions could be refined, and even more criteria could be added. However, this list gives a good impression of the diversity of a management problem analysis.

2.1 Polling Frequency

Management applications may have different needs for the polling frequency on managed objects. Reasons for high polling frequencies might be the demand for monitoring or accounting at a high time resolution or the lack of asynchronous notifications although timely handling of error conditions is required.

2.2 Number of Agents

Obviously, the number of agents has direct impact on the network management traffic that is induced to the network.

2.3 Number of Managed Objects

Similar to the number of agents that are involved with a single management task, the number of managed objects influences the network traffic. Only to a limited degree, this traffic can be reduced by gathering multiple variables in common protocol requests.

2.4 Network Topology

When a large physical network has a hierarchical topology, e.g., in a global company with multiple offices, it might be reasonable to follow the network topology when assigning management tasks to network elements.

2.5 Administrative Domains

Compared to the network topology there might be separate administrative domains in a network, e.g., the management of a network application service might be separated from the network connectivity management. To keep these administrative domains independent, it might be reasonable to separate also their management tasks with respect to their locations and processing environments.

2.6 Security Considerations

Similar to the administrative reasons it could required to separate different management tasks because of their different security requirements. E.g., the management of employee information and their passwords usually has higher security requirements than the monitoring of network equipment.

2.7 Reliability Requirements

The reliability of management functions depends on the reliability of the underlying network. E.g., error conditions in a LAN cannot be recognized and eliminated if an unreliable WAN connection separates the responsible manager from the affected LAN. Thus, highly reliable manager actions cannot be achieved in this case.

2.8 Network Performance

Similar to the reliability of the underlying network, its throughput and delay might also be considered, taking the average and maximum network load into account.

3 Characteristics of Management Systems

In this section, some management system characteristics are presented that are significant to find a suitable management distribution model. Each characteristic is described by its dependencies on the aforementioned criteria.

3.1 Numbers of Managers and Elements

If the amount of management traffic between a top-level manager and the agents controlled by this manager is relatively high in comparison to the network bandwidth, a decentralized paradigm becomes more suitable. Obvious reasons for higher management traffic are high polling frequencies, large numbers of agents controlled by a single manager, and huge amounts of communicated management information, probably due to semantically rich interactions [4].
As a consequence the number of managers m (top-level manager and mid-level managers) can be increased in relation to the total number of elements n in the system, so that the CPU and network load can be distributed among these managers.
Furthermore, all criteria that increase the heterogeneity of the management system (see Heterogeneity, section 3.3), obviously also increase the number of managers m.

3.2 Connectivity between Elements

Besides the number of elements in a management system, the number of inter-manager communication relations c divided by the number of managers m can be regarded as a significant characteristic, that we call the *connectivity* $C = c/m$.
Reasons for a higher connectivity might be provisions to increase redundancy for better reliability, or a fine grained distribution of communicating sub-tasks of a single management task.

3.3 Heterogeneity

Security considerations, various administrative domains, and reliability requirements can lead to a number of separate management tasks. Hence, these criteria can increase the heterogeneity of the management system as a whole.

3.4 Delegation Domains

The decision which tasks are delegated to which managers can be biased by different criteria. In cases of wide distributed system elements and unreliable or slow links, geographical or topological reasons might be significant, so that it is reasonable to locate mid-level managers near the corresponding agents. In other situations, it might be meaningful to separate functional blocks (see Heterogeneity, section 3.3) and build up delegation domains by different functional management tasks, may be also on tightly located mid-level managers.

3.5 Complexity of Management Tasks

Another criterion is the distribution of the complexity of management tasks among the managers of a management system. The higher the complexity of a manager at level N the more exonerated are the parent managers at levels $< N$ and the network links to them.
The criteria for increasing the complexity of mid-level managers are similar to those for splitting delegation domains (section 3.4).

3.6 Dynamic Delegation

While traditional management tasks are fulfilled by static code at a single static location, distributed management systems may show different degrees of dynamic task assignment and relocation. This paradigm is often called *mobile code*. It can be quantified by the number of rules that cause changes to the managers' task topology and the frequency of topology changes.
The degree of dynamic delegation that is adequate for a given scenario depends on the level of abstraction used to model a management application [3]. E.g., in a management system where a manager has to monitor all of the network's routers without knowing them individually, it might be suitable to locate them dynamically in order to delegate management tasks to every currently detected subnet. On the other hand, if the manager is configured so that it statically knows all subnets, dynamic delegation is not required.

3.7 Process Mobility

While dynamic delegation addresses only the mobility of static code, mobile processes affect also the delegation of running management processes including their execution states. This paradigm is often called *mobile agents*. Similar to dynamic delegation, process mobility can be quantified by the frequency and number of rules that make processes migrate from one mid-level manager to another. Mobile agents might be suitable for management systems which are described through the use of very high-level abstractions [3].

4 Management Distribution Models

In the first section four basic management distribution models have been enumerated. This sections shows how combinations of characteristics presented in section 3 can be mapped to these basic models.
The typical *centralized management* architecture consists of two levels: a single (top-level) manager and a number of agents. In this case, the connectivity (section 3.2) is zero. In contrast, if the number of managers m is approximately the total number of elements n, and the connectivity C is high, we talk about *cooperative management*. In between, we can differ two other forms: *weakly distributed management* and *strongly distributed management*. Figure 2 clarifies this classification. Figure 3 shows the structures of the distribution models.

$1 = m$	$C = 0$	centralized management
$1 < m \ll n$	$0 < C < 1$	weakly distributed management
$1 \ll m < n$	$1 \leq C \ll m$	strongly distributed management
$m \approx n$	$1 \ll C \leq m$	cooperative management

Figure 2: Classification by numbers of managers and elements.

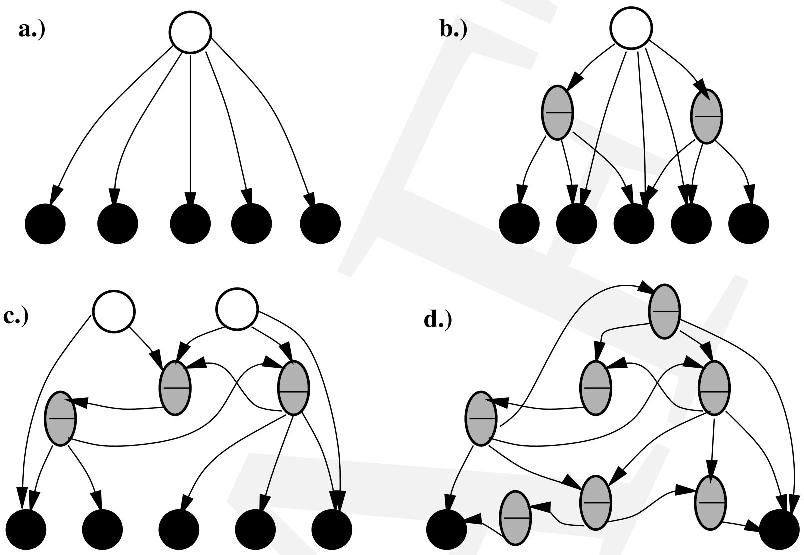

Figure 3: Structures of different classes of management systems. An open circle denotes a top-level manager, a grey split ellipse a mid-level manager, and a filled circle an agent. **a.)** centralized, **b.)** weakly distributed, **c.)** strongly distributed, **d.)** cooperative management system.

The numbers of managers and elements (section 3.1) and their connectivity (section 3.2) make up a management system's topology. Typically, centralized and weakly distributed management systems show up a strictly hierarchical *vertical delegation* of management tasks, i.e. (assuming that the top-level manager is at level 1) managers at level N delegate tasks only to managers at levels $> N$. In contrast, in strongly distributed and cooperative management the mid-level managers establish communication relationships between each other, which implies *horizontal delegation*, so that there is no clean hierarchy.
Heterogeneous management systems (section 3.3) require at least a weakly distributed management architecture. However, the degree of heterogeneity does not interfere with the degree of more complex distribution models, like strongly distributed or cooperative management. Similar to heterogeneous management systems the separation of multiple management tasks by domains and with some kind of distributed complexity (sections 3.4 and 3.5) implies at least a weakly distributed management architecture. Typically, the complexity of tasks is more balanced in strongly distributed and cooperative management systems.

Dynamic delegation (section 3.6) of management tasks is already possible in weakly distributed management systems, but more convenient in strongly distributed and cooperative management systems. Efficient process mobility (section 3.7) can only be achieved in management systems that allow horizontal delegation. Hence, it looks most convenient to apply mobile agents to cooperative management systems.

5 Conclusion

To find a suitable distributed management model a systematic analysis of the given management problem is essential. This article presents a number of management problem criteria and rules that lead to typical characteristics of reasonable management systems, which can in turn be mapped to four basic management distribution models. The criteria, rules and characteristics can be refined and completed in future. However, the presented process gives an impression of a systematic management problem analysis.

References

[1] Y. Yemini, G. Goldszmidt, and S. Yemini. Network Management by Delegation. In *Proc. International Symposium on Integrated Network Management*, pages 95–107, April 1991.

[2] J. Schönwälder, J. Quittek, and C. Kappler. Building Distributed Management Applications with the IETF Script MIB. *IEEE Journal on Selected Areas of Communications*, 18(5), May 2000.

[3] J. P. Martin-Flatin, S. Znaty, and J. P. Hubaux. A Survey of Distributed Enterprise Network and Systems Management Paradigms. *Journal of Network and Systems Management*, 7(1):9–26, March 1999.

[4] K. Meyer, M. Erlinger, J. Betser, C. Sunshine, G. Goldszmidt, and Y. Yemini. Decentralizing Control and Intelligence in Network Management. In *Proc. 4th International Symposium on Integrated Network Management*, May 1995.