The present thesis work introduces some recent and novel results about the problem of fault diagnosis for distributed nonlinear and large scale systems. The problem of automated fault diagnosis and accommodation is motivated by the need to develop more autonomous and intelligent systems that operate reliably in the presence of system faults. In dynamical systems, faults are characterized by critical and unpredictable changes in the system dynamics, thus requiring the design of suitable fault diagnosis schemes. A fault diagnosis scheme that drew considerable attention and provided remarkable results is the so called *model based* scheme, which is based upon a mathematical model of the healthy behaviour of the system that is being monitored. At each time instant, the model is used to compute an estimate of what should be the current behaviour of the system, assuming it is not affected by a fault. If the behaviour of the system is characterized by the time evolution of its *state* vector $x(t)$, and the *inputs* to the system are denoted as $u(t)$, then the most general nonlinear and uncertain discrete time model can be represented by

$$x(t+1) = f(x(t), u(t)) + \eta(t)\,,$$

where the nonlinear function f represents the nominal model of the healthy system, and $\eta(t)$ is an uncertainty term. A proven way to compute an estimate $\hat{x}(t)$ of the state $x(t)$ is by using the following estimator

$$\hat{x}(t+1) = \lambda(\hat{x}(t) - x(t)) + f(x(t), u(t))\,,$$

so that in healthy conditions the *residual* $\epsilon(t) \triangleq x(t) - \hat{x}(t)$ is, in practice, close to zero. Should the residual cross at a certain point a suitable *threshold* $\bar{\epsilon}(t)$, the observed difference between the model estimate and the actual measurements will be explained by the presence of a fault.

The model-based scheme outlined so far has showed many interesting properties and advantages over signal-based ones, but anyway poses practical implementation problems when one tries to apply it to actual distributed, large-scale systems. In fact an implicit assumption about the model-based scheme is that the task of measuring all the state and input vectors components, and the task of computing the estimate $\hat{x}(t)$ can be done in real-time by some single and powerful computer. But for large enough systems, this assumptions cannot be fulfilled by available measurement, communication and computation hardware. This problem contitutes the motivation of the present work. It will be solved by developing decomposition strategies in order to break down the original centralized diagnosis problem into many distributed diagnosis subproblems, that are tackled by agents called *Local Fault Diagnosers* that have a limited view about the system, but that are allowed to communicate between neighbouring agents. In order to take advantage of the distributed nature of the proposed schemes, the agents are allowed to cooperate on the diagnosis of parts of the system shared by more than one diagnoser, by using consensus techniques.

Chapter 2 introduces the problem of model-based fault diagnosis by presenting recent results about the centralized diagnosis of uncertain nonlinear discrete time systems. The development of a distributed fault diagnosis architecture is covered in the key Chapter 3, while Chapters 4 and 5 show how this distributed architecture is implemented for discrete and continuous time nonlinear and uncertain large–scale systems. In every chapter an illustrative example is provided, as well as analytical results that characterize the performances attainable by the proposed architecture.