

Abstract. We use the failures model of CSP to describe the behaviour of a class of networks of communicating processes. This model is well suited to reasoning about the deadlock potential of networks. We introduce a number of simple conditions on networks which aid deadlock analysis either by localizing the analysis required for a proof of deadlock-freedom or by restricting the circumstances in which deadlock could occur. In particular, we formulate some simple theorems which characterize the states in which deadlock can occur, and use them to prove some theorems on the absence of global deadlock in systems. We identify a special class of unidirectional networks and develop specialized results on their deadlock-freedom. We develop more general methods based on (at most) pairwise local deadlock analysis in networks, applicable to the large class of conflict-free networks. We introduce a methodology for proving deadlock-freedom in a large network by decomposing it into subnetworks which can be analysed separately. A variety of examples is given to show the utility of these results. We compare our work with earlier work by several other authors, and make some suggestions for future research.

Key words: Deadlock – Networks – Communicating processes