C. M. Ozveren, A. S. Willsky and P. J. Antsaklis, "Stability and Stabilizabilty of Discrete Event Dynamic Systems," Journal of the Association of Computing Machinery, Vol 38, No 3, pp 730-752, 1991.

Journal of ACM

February 1989

LIDS-P-1853

STABILITY AND STABILIZABILITY OF DISCRETE EVENT DYNAMIC SYSTEMS[1]

Cüneyt M. Özveren[2]

Alan S. Willsky[2]

Panos J. Antsaklis[3]

February 14, 1989

Abstract

A finite state automaton is adopted as a model for Discrete Event Dynamic Systems (DEDS). Stability is defined as visiting a given set E infinitely often. Stabilizability is defined as choosing state feedback such that the closed loop system is stable. These notions are proposed as properties of resiliency or error-recovery. An important ingredient in stability is shown to be a notion of transiton-function-invariance. Relations between our notions of stability and invariance, and the notions of safety, faireness, livelock, deadlock, etc. in computer science literature are pointed out. Connections are estabished between our notions of invariance and the classical notions of A-invariance and (A, B)-invariance of linear systems. Polynomial algorithms for testing stability and stabilizability, and for constructing a stabilizing control law are also presented.

[1]Research supported by the Air Force Office of Scientific Research under Grant AFOSR-88-0032 and by the Army Research Office under Grant DAAL03-86-K0171. The research of the first two authors was partially done during their stay at Institut de Recherche en Informatique et Systèmes Aléatoires (IRISA), Rennes, France, and the second author was also supported by IRISA during this time.

[2]Laboratory for Information and Decision Systems, MIT, Cambridge, MA 02139.

[3]University of Notre Dame, Notre Dame, Indiana, 46556.

Journal of the Association of Computing Machinery, Vol 38, No 3, pp 730-752, 1991.

Index Terms

Automata Theory
Control Theory
Discrete Event Dynamic Systems
Flexible Manufacturing Systems
Resiliency
Stability

1 Introduction

Discrete Event Dynamic Systems (DEDS) have received considerable attention in the control literature recently. Many large scale dynamic systems seem to have a DEDS structure, at least at some level of description. Some examples are manufacturing systems [12,18], communication systems (such as data networks, and distributed systems) [3], and expert systems (such as CPU design, or air-traffic management) [7,9,21].

The notion of the control of a DEDS was, to our knowledge, first explicitly introduced in the work of Wonham, Ramadge, et al. [10,13,15,16,22,23]. In this work, it is assumed that certain events in the system can be enabled or disabled. The control of the system is achieved by choice of control inputs that enable or disable these events. The objective is to have a closed loop system, so that the event trajectory in this system is always in a given set of desired strings of events. This approach is generally classified as a linguistic approach, since the objective is defined in terms of the language generated by the closed-loop system, i.e., the set of possible strings of events.

This work has prompted a considerable response by other researchers in the field, and one of the principal characteristics of this research has been the exploration of alternate formulations and paradigms that provide the opportunity for new and important developments building on the foundations of both computer science and control. The work presented here is very much in that spirit with, perhaps, closer ties to more standard control concepts. In particular, in our work, we have had in mind the development of the elements needed for a regulator theory for DEDS. In this paper, we develop notions of stability and stabilizability for DEDS which might,

more concretely, be thought of as properties of resiliency or error-recovery.

The goal in the work of Wonham, Ramadge, et al. is to restrict the behavior of the system so that all strings generated from the given initial state are in a given set of "legal" strings. In a sense, what we seek here is to develop control methods for re-establishing such legal behavior following the occurence of one or more anomalous events. For example, a manufacturing system is always subject to failures. Thus, in the Wonham and Ramadge context, one would include all possible strings with failures and successful recoveries in the legal language (the set of legal strings). In our formulation, we focus on states rather than strings. Specifically, assume that we have identified the set of "good" states, i.e., the set of initial states, from which only legal strings are generated. Our focus then is to test if all trajectories from other states visit the "good" states infinitely often, so that the system recovers from any possible error in a finite number of transitions. If, in fact, failures do not happen very often, then the system will exhibit legal behavior "most of the time".

Another goal of this work is to establish connections with the related notions in computer science. In particular, the concept of stability we use here has been introduced by researchers in a number of different computer science contexts. What distinguishes our work and makes it of potential interest in computer science as well as in control theory is the introduction of control and feedback to formulate and solve the problem of stabilizing systems. For example, our notion of pre-stability and the algorithm we provide is exactly the same as the notion of inevitable reachability and the algorithm of Sifakis [17]. Other notions of Sifakis can be characterized using stability and transition-function-invariance (f-invariance of Section 2). Thus, our results in stabilizability can be directly applied to his notions if control was to be introduced

in his framework. In [6], a system is defined to be self-stabilizing if starting at any unsafe state, it is guaranteed to reach a "safe" state within a finite number of transitions. This is also the same as our notion of stability, and therefore our results of stabilizability can be applied to this case. Finally, fair execution sequences, in concurrent systems, are defined as those execution sequences in which each process is executed infinitely often, [4]. This notion is also connected to our notion of stability and our results on stabilizability can tell us how to achieve fairness.

In the next section, we introduce the mathematical framework considered in this paper and address the problem of stability. In particular, we first introduce a notion of pre-stability, which is based on testing if all trajectories from a state go through the "good" states. We then define transition-function-invariance (f-invariance) in our framework and characterize stability in terms of pre-stability and f-invariance. We also present algorithms for testing pre-stability and stability. In Section 3, we address the problem of stabilizability. We first present an algorithm to construct a pre-stabilizing state feedback, and we classify different pre-stabilizing feedbacks by the degree of restriction imposed on the behavior. As an extension of f-invariance, we examine achieving f-invariance by control inputs and we then impose the constraint that f-invariance is achieved while keeping the system alive. We combine this with pre-stabilizability and present an algorithm to construct a stabilizing state feedback. Finally, in Section 4, we summarize our results and outline further research directions.

2 Stability

In this section, we define our notion of stability and provide an algorithm that tests stability.

2.1 Preliminaries

The class of systems we consider are nondeterministic finite-state automata. The basic object of interest is:

$$G = (X, \Sigma) \tag{2.1}$$

where X is the finite set of states, with $n = |X|$, and Σ is the finite set of possible events. The dynamics of the system are characterized by two functions f and d:

$$x[k+1] \in f(x[k], \sigma[k+1]) \tag{2.2}$$

$$\sigma[k+1] \in d(x[k]) \tag{2.3}$$

Here $x[k] \in X$ is the state after the kth event, and $\sigma[k+1] \in \Sigma$ is the $(k+1)$st event. The function $d : X \to 2^{\Sigma}$ is a set valued function that specifies the set of possible events defined at each state (so that, in general, not all events are possible from each state), and the function $f : X \times \Sigma \to X$ is also set valued, so that the state following a particular event is not necessarily known with certainty. The triple $A = (G, f, d)$ representing our system can also be visualized graphically as in Figure 2.1. Here circles denote states, and arcs denote transitions, where the symbol in each arc label denotes the event corresponding to that transition. Thus, in this example, $X = \{0, 1, 2, 3\}$, $\Sigma = \{\alpha, \beta, \delta\}$, and, for example, $d(1) = \{\alpha, \delta\}$, $f(0, \beta) = \{0, 3\}$, etc. A transition, also denoted as $x \to^{\sigma} y$, consists of a source state, x, an event, $\sigma \in d(x)$, and a destination state, $y \in f(x, \sigma)$.

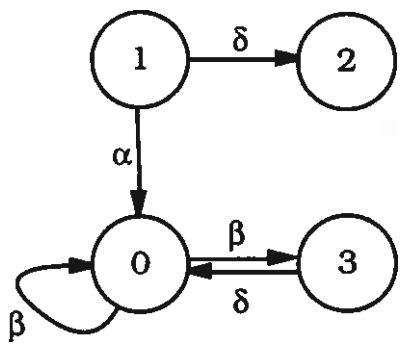

Figure 2.1: Simple Example

In the subsequent sections, we will use the following terminology concerning state and event trajectories:

- A finite string of states, $\mathbf{x} = x_0 x_1 \cdots x_j$ is termed a path or a state trajectory from x_0 if $x_{i+1} \in f(x_i, d(x_i))$ for all $i = 0 \cdots j-1$, where

$$f(x, d(x)) = \bigcup_{\sigma \in d(x)} f(x, \sigma)$$

  We say $x \in \mathbf{x}$ if $x_i = x$ for some i. Let $\mathcal{X}(A, x)$ denote the set of all possible paths from x. A path is termed a cycle if $x_0 = x_j$ and a cycle is termed a primary cycle if there exists no distinct pair $i_1, i_2 \in 0 \cdots j-1$ such that $x_{i_1} = x_{i_2}$, i.e., if it contains no other cycles. For example, in Figure 2.1, 12, 3003, and 030 are all paths, 3003, and 030 are cycles, and 030 is a primary cycle. In general, there may be infinitely many cycles, but only a finite number of primary cycles. For example, the primary cycles of Figure 2.1 are 00, 030, and 303.

- Similarly, a finite string of events $s = \sigma_1 \cdots \sigma_j$ is termed an event trajectory from $x \in X$ if $\sigma_1 \in d(x)$ and $\sigma_{i+1} \in d(f(x, \sigma_1 \cdots \sigma_i))$ for all i, where we extend

f to Σ^* via

$$f(x, \sigma_1 \cdots \sigma_i) = f(f(x, \sigma_1 \cdots \sigma_{i-1}), \sigma_i)$$

with $f(x, \epsilon) = x$, where ϵ is the empty string. In Figure 2.1, $\alpha\beta\beta\delta$ is an event trajectory. Let $L(A, x)$ denote the set of all possible event trajectories starting from state x.

For most applications of interest to us, it is desired that the system can never reach a point at which no event is possible. This is a notion of liveness. For example, a manufacturing system should, always, be capable of producing something. We use

$$R(A, x) = \{y \in X | x \rightarrow^* y\}$$

to represent the set of states that can be reached from x, where $\rightarrow^*$ denotes any number of transitions, including no transitions. Thus, $R(A, x)$ always includes x. For example, in Figure 2.1, the reach of 1 is X itself. Intuitively, we define a state to be alive if all event trajectories from that state have infinite length extensions. We let

$$D = \{x \in X | d(x) = \emptyset\} \tag{2.4}$$

denote the set of states which have no events defined, and term these the <u>dead</u> states. We formally define liveness as follows:

Definition 2.1 A state $x \in X$ is <u>alive</u> if $d(y) \neq \emptyset$ for all $y \in R(A, x)$. A subset Y of X is termed a live set if all $x \in Y$ are alive. A system A is termed alive if X is a live set. ❑

For example, in Figure 2.1, $D = \{2\}$ and 0 is alive, whereas 1 is not. Clearly, the class of live sets is closed under arbitrary unions and intersections. The maximal live subset of X, X_a, is given by the set of states that <u>cannot</u> reach the dead states:

$$X_a = \overline{R(A^{-1}, D)} \tag{2.5}$$

where overline denotes the complement and A^{-1} denotes A with the transitions reversed, i.e., $A^{-1} = (G, f^{-1}, d^{-1})$ where:

$$f^{-1}(x,\sigma) = \{y \in X | x \in f(y,\sigma)\} \tag{2.6}$$

$$d^{-1}(x) = \{\sigma \in \Sigma | \exists y \in X \text{ such that } x \in f(y,\sigma)\} \tag{2.7}$$

Thus, $R(A^{-1}, D)$ is the set of states that can reach D in the original system. For example, in Figure 2.1, $D = \{2\}$ and $X_a = \{0,3\}$. Note that the state 1 is not alive since there exists a trajectory from 1 which goes to state 2, which is a dead state.

Note that if we replace D by any set of states, then X_a is the maximal set of states that avoid D. This situation arises, for example, in mutual exclusion problems, in the context of manufacturing systems, [11], and computer systems, [1,2,5], where a number of users compete for a limited number of resources, say r. In that case, the states that represent $p > r$ users attempting to use the resources are undesirable and one wants to find the set of states that avoid this violation. Golazewski and Ramadge, [8], address this problem, in the context of DEDS problems that consist of many interacting components. Our contributions to this problem are presented in [14].

We conclude this section by presenting an algorithm to compute the reach of a set of states. It immediately follows from the definition that the reach of a set of states, X_0, is the fixed point of $R = f(R, \Sigma) = \bigcup_{x \in R} f(x, d(x))$ such that $R \supset X_0$. Thus, we have the following algorithm:

Proposition 2.2 The following algorithm computes $R(A, X_0)$ for any $X_0 \subset X$ and it has complexity $O(n)$:

Algorithm Let $R_0 = Q_0 = X_0$ and iterate:

$$R_{k+1} = R_k \cup f(Q_k, \Sigma)$$

$$Q_{k+1} = R_{k+1} \cap \overline{R_k}$$

Terminate when $R_{k+1} = R_k$.

Proof: Clearly, the algorithm terminates in a finite number of steps, say r, and $R_r = R(A, X_0)$. Since each state is visited only once, the complexity of the algortihm is $O(n)$. ❑

For example, in Figure 2.1, in order to compute the reach of 1, we have: $R_0 = \{1\}, R_1 = \{0, 1, 2\}, R_2 = X, R_3 = X$, and the algorithm terminates. Thus, the reach of 1 is X.

2.2 Pre-Stability

The notion of stability we wish to capture can be thought of as error recovery. Specifically, as in Wonham and Ramadge, one can imagine a set of desired event trajectories that one would like to see in a DEDS. For example, in a manufacturing system, these sequences might consist of a concatenation of subsequences each of which corresponds to the successful production of an individual component and the return of the system to a "start-up" state, from which it can initiate the next production task. Because of the possibility of failures or errors, actual behavior may deviate from this ideal, and what one would like is that after such a failure, the system recovers. To capture this idea, we suppose that we have identified a subset E, of the state space X, so that returning to E corresponds to being in a position to continue desired behavior from that point on. For example, in a manufacturing system, E might be the set of all non-failure states or it might simply be the set of start-up states, which, in the Wonham and Ramadge framework, can be thought of as the initial state from which legal event trajectories are generated. Error recovery or stability then corresponds to

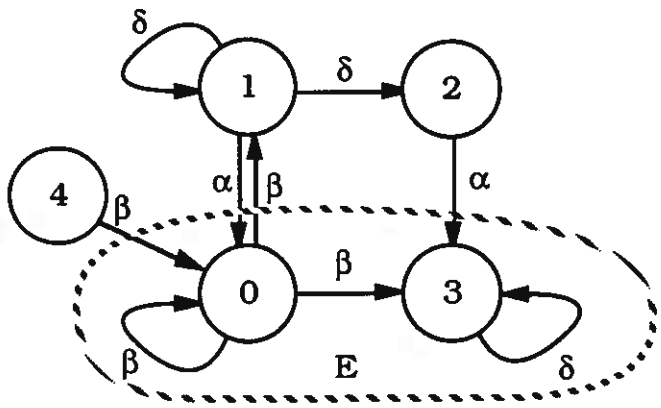

Figure 2.2: Stability Example

a return to E in a finite number of transitions following any excursion out of E. There is a useful linguistic interpretation of this concept. Define the desired language as

$$L(A, E) = \bigcup_{x \in E} L(A, x) \tag{2.8}$$

What we would like is the following. Suppose that for some reason, we are in a state $x \notin E$. Then, we want all possible event trajectories from x to differ from a desired trajectory by at most a finite prefix.

Given E, we define a state $x \in X$ to be stable if all paths from x go through E in a finite number of transitions and then visit E infinitely often. For example, in Figure 2.2, where $E = \{0, 2\}$, only 2 and 3 are stable states. State 1 is not stable since the system can loop at 1 for an infinite number of transitions. States 0 and 4, although 0 is in E, are not stable either since the sytem can make a transition to state 1 and stay there forever. This notion is similar to but not exactly the same as the notion of Büchi acceptance [20]: An infinite event trajectory is Büchi acceptable if there exists a corresponding state trajectory which visits a terminal state (where a

set of terminal states is given) infinitely often. In our notion of stability, if we let E be the set of terminal states, all possible state trajectories, from a stable state, visit E infinitely often.

Our notion of stability is captured in two stages. We term x pre-stable if all paths from x go to E in a finite number of transitions. In other words, no path from x ends up in a cycle that does not go through E. For example, in Figure 2.2, 0, 2, 3, and 4 are pre-stable. This notion is exactly the same as the notion of inevitable-reachability of Sifakis [17]. A state is then stable if all the states in its reach are pre-stable. In Figure 2.2, 0 and 4 are not stable since they can reach 1, which is not pre-stable.

We formalize pre-stability as follows:

Definition 2.3 Given a live system A and some $E \subset X$, a state $x \in X$ is pre-stable with respect to E if for all $\mathbf{x} \in \mathcal{X}(A, x)$ such that $|\mathbf{x}| \geq n$, there exists $y \in \mathbf{x}$ such that $y \in E$. ❑

We say that a set of states is pre-stable with respect to E if all its elements are pre-stable and a system A is pre-stable with respect to E, if all states in X are pre-stable.

If all paths from x go through E, then they do so in less than n transitions since otherwise x has a a cycle that does not go through E, and thus, there exists a path which never goes through E. Equivalently, pre-stability can be characterized in terms of the primary cycles. To formalize this, let $R_E(A, x)$ denote the set of states that can be reached by trajectories from x that do not go out of E if they enter E at all. For example, in 2.2, $R_E(A, 4) = \{0, 3, 4\}$. In other words, $R_E(A, x) = R(A', x)$ where A' is an automaton created from A by removing all transitions, from states in E, which take that state outside of E. For example, in Figure 2.2, we only remove the

transition $0 \rightarrow^{\beta} 1$. Then, x is pre-stable if and only if all primary cycles in $R_E(A, x)$ go through E. In Figure 2.2, the self-loop at 1 is a primary cycle and does not go through E.

Proposition 2.4 Given a live A and $x \in X$, x is pre-stable with respect to E iff all primary cycles in $R_E(A, x)$ include at least one state in E. In general, A is pre-stable with respect to E iff all primary cycles in X include at least one state in E.

Proof: Straightforward by assuming the contrary in each direction. ❑

The class of sets, that are pre-stable with respect to E, is closed under arbitrary unions and intersections. Now, we derive an algorithm that computes X_P, the maximal subset of X that is pre-stable with respect to E. Our algorithm, which is the same as the one given for inevitable reachability in Sifakis, is based on starting from E and growing the currently known set of pre-stable states by including, at each step, those states x such that $f(x, d(x))$ is a subset of the current set.

In developing this algorithm, we first need the following lemma, which states that if some state x is pre-stable, then either x is in E or all the events defined from x take x to a pre-stable state:

Lemma 2.5 $x \in X$ is pre-stable with respect to E iff $x \in E$ or $f(x, d(x))$ is pre-stable with respect to E.

Proof: Straightforward. ❑

Also, note that given a set of pre-stable states Q, that include E, we can test pre-stability of other states by testing pre-stability with respect to Q:

Lemma 2.6 Given $Q_1, Q_2 \subset X$ such that Q_1 is pre-stable with respect to E and $Q_1 \supset E$, Q_2 is pre-stable with respect to E iff Q_2 is also pre-stable with respect to Q_1.

Proof: ($\rightarrow$) Obvious since $Q_1 \supset E$.

($\leftarrow$) Suppose that some $x_2 \in Q_2$ goes to a cycle that avoids E. By Proposition 2.4, this cycle goes through some state $x_1 \in Q_1$. Then, x_1 cannot be pre-stable with respect to E, and we have a contradiction. ❑

These two lemmas lead to the following algorithm:

Proposition 2.7 The following algorithm computes X_P, and it has complexity $O(n^2)$:

Algorithm Let $X_0 = E$ and iterate:

$$X_{k+1} = \{x | f(x, d(x)) \subset X_k\} \cup X_k$$

Terminate when $X_{k+1} = X_k$.

Proof: Clearly, X_0 is pre-stable with respect to E. Suppose that X_k is pre-stable with respect to E. By Lemma 2.5, X_{k+1} is pre-stable with respect to X_k and by Lemma 2.6, it is also pre-stable with respect to E. This algorithm terminates in at most n steps. Let us say that it terminates in r steps. Suppose that there exists some $x_1 \in X_P$ such that $x_1 \notin X_r$, then there exits $\sigma \in d(x_1)$ and $x_2 \notin X_r$ such that $x_2 \in f(x_1, \sigma)$. The same holds true for x_2 and some $x_3 \notin X_r$, etc. Thus, there exists a path which never reaches X_r. Since also $X_r \supset E$, x_1 is <u>not</u> pre-stable with respect to E and we have a contradiction. Finally, to justify complexity, note that this algorithm terminates in at most n iterations. Since all states can be visited at most once at each iteration, the complexity of the algorithm is $O(n^2)$. ❑

In Figure 2.2, $X_1 = X_2 = X_P = \{0, 2, 3, 4\}$. Note that the number of steps in which this algorithm terminates, is a notion of radius for the pre-stable part of the system where E is taken as the center. That is, it is precisely the length of the maximum length trajectory between any pre-stable state and the state in E at which

this trajectory enters E for the first time. In Figure 2.2, this radius is one. In fact, if some x is included at step k of the algorithm, then the maximum number of transitons it takes to go from x to E is k.

2.3 Stability and f-Invariance

As motivated in the previous sections, we define stability as follows:

Definition 2.8 Given a live system A and some $E \subset X$, a state $x \in X$ is <u>stable</u> with respect to E. if all infinite state trajectories starting from x visit E infinitely often. More precisely, x is stable if for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}(A, x)$ so that $\mathbf{x}_2 = \mathbf{x}_1\mathbf{z}$, with $|\mathbf{z}| \geq n$, then there exists $y \in \mathbf{z}$ such that $y \in E$. ❑

This definition states that at any point in a trajectory from a stable state, we know that the trajectory will re-visit E within a finite number of transitions. In Figure 2.2, clearly, 1 is not stable. States 4 and 0 are not stable because there exist trajectories that start from these states and go to state 1, and subsequently, these trajectories may loop in state 1 forever.

An immediate consequence of this definition is the following, which states that the stability of a state is equivalent to the pre-stability of its reach.

Proposition 2.9 Given a live A and $x \in X$, x is stable with respect to E iff $R(A, x)$ is pre-stable with respect to E. ❑

A subset of X is stable if all its elements are stable, and a system A is termed stable if X is a stable set. We immediately have:

Proposition 2.10 A is a system stable with respect to E iff it is also pre-stable with respect to E. ❑

We also have the following counterpart of Proposition 2.4:

Proposition 2.11 Given a live A and $x \in X$, x is stable with respect to E iff all primary cycles in $R(A, x)$ include at least one state in E. In general, A is stable with respect to E iff all primary cycles in X include at least one state in E.

Proof: Straightforward. ❑

If we compare this to Proposition 2.4, note that the second statements are exactly the same. This is due to Proposition 2.10.

The class of sets stable with respect to E is closed under arbitrary unions and intersections. Let X_S denote the maximal set stable with respect to E (Note that X_S can be the empty set, for example, if we let $E = \{0\}$ in Figure 2.2). Then, X_S is the set of states in X_P from which we can only reach pre-stable states. Let us first formalize this notion of staying within a given set of states (corresponding to the notion of A-invariant subspaces of system thoery):

Definition 2.12 A subset Q of X is <u>f-invariant</u> if $f(Q, d) \subset Q$ where

$$f(Q, d) = \bigcup_{x \in Q} f(x, d(x))$$

❑

It immediately follows that any trajectory that starts in an f-invariant set stays in that set:

Proposition 2.13 Q is f-invariant iff $R(A, Q) \subset Q$.

Proof: Straightforward. ❑

The class of f-invariant sets is closed under arbitrary unions and intersections. We then have:

Proposition 2.14 X_S is the maximal f-invariant set in X_P.

Proof: ($\subset$) Clearly, $X_S \subset X_P$. Also, X_S is f-invariant since if a state $x \in X_S$ can reach a state that is <u>not</u> stable, then x cannot be stable.

($\supset$) Let X_f denote any f-invariant set in X_P. Any path from a state in X_f goes through E, and if gets out of E it stays in X_f and thus in X_P. Therefore, X_f is stable. ❑

Note that the maximal f-invariant set in Q is the set of states in Q that <u>cannot</u> reach any state in $\overline{Q}$, i.e., it is $\overline{R(A^{-1}, \overline{Q})}$. Thus, we can compute X_S as follows:

$$X_S = \overline{R(A^{-1}, \overline{X_P})} \tag{2.9}$$

3 Stabilizability

So far, we have dealt with notions that are close to those commonly seen in the automata theory literature. In this section, we introduce control and reconsider the notions formulated in the previous section. We define pre-stabilizability (respectively stabilizability) as finding a state feedback such that the closed loop system is pre-stable (stable). We present a polynomial algorithm for constructing a pre-stabilizing feedback. This algorithm is a natural extension of Algorithm 2.7 and it generates a state feedback which is maximally restrictive, in the sense that it disables as many events as possible at each state, and path minimizing, in the sense that the maximum length path from pre-stable states to E is minimized. We also present an algorithm for constructing a minimally restrictive feedback. Finally, we introduce a notion of (f, u)-invariance, achieving f-invariance by choice of state feedback, and use this notion, together with the constraint that the closed loop system needs to be alive, to develop a polynomial algorithm for constructing a stabilizing feedback.

3.1 Pre-Stabilizability

To introduce control, we modify our system model as follows:

$$G = (X, \Sigma, U) \tag{3.1}$$

where, U is the set of admissible control inputs. We introduce control by allowing certain events at each state to be disabled. This framework is the same as that of Ramadge and Wonham, except that in our case, an event that is controllable at one state may not be controllable at another state. We let $U = 2^{\Sigma}$ and the dynamics are

described by:

$$x[k+1] \in f(x[k], \sigma[k+1]) \tag{3.2}$$

$$\sigma[k+1] \in (d(x[k]) \cap u[k]) \cup e(x[k]) \tag{3.3}$$

where, $u[k] \in U$ is the control input after the kth event, and $e : X \to 2^{\Sigma}$ is a set valued function that specifies the set of events that cannot be disabled at each state. Without loss of generality, we assume that $e(x) \subset d(x)$ for all x. The interpretation of 3.3 is straightforward. The set $d(x)$ represents an "upper bound" on the set of events that can occur at state x—no matter what we do, the next event will be in $d(x)$. The set $e(x)$, on the other hand, is a lower bound—no matter what we do, any event in $e(x)$ may still occur. The effect of our control action is adjusting the set of possible events between these bounds, by disabling some of the controllable events, i.e., elements of the set $d(x) \cap \overline{e(x)}$. Note, therefore that while in our definition $u(x)$ can be any subset of Σ, we can, without loss of generality assume that $u(x) \subset d(x) \cap \overline{e(x)}$. We will make this assumption throughout. The quadruple $A = (G, f, d, e)$ representing our system can also be visualized graphically as in Figure 3.1. In addition to the graphical representation of the previous section, we mark the controllable events by ":u". For example, in Figure 3.1, γ is controllable at 1 and 2, whereas β is controllable only at 3.

A state feedback law is a map $K : X \to U$. Given a state feedback K, let $A_K = (G, f, d_K, e)$ denote the closed loop system where

$$d_K(x) = (d(x) \cap K(x)) \cup e(x) \tag{3.4}$$

We define pre-stabilizability as follows:

Definition 3.1 Given a live system A and some $E \subset X$, $x \in X$ is pre-stabilizable with

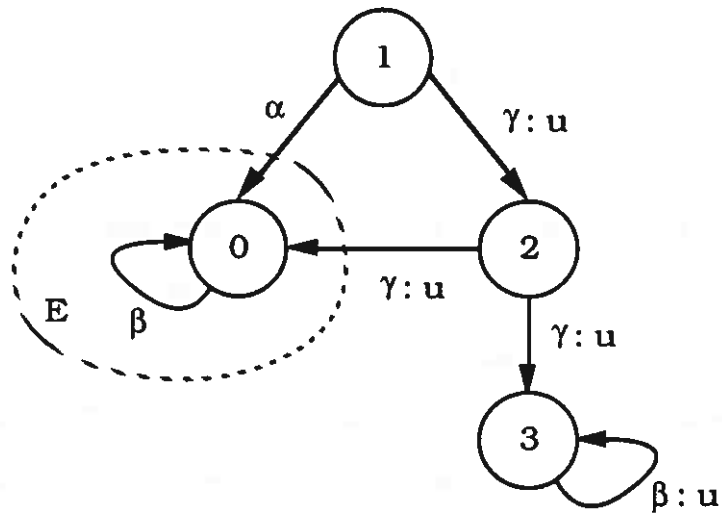

Figure 3.1: Example for the Notion of Pre-Stabilizability

respect to E if there exits a state feedback K such that x is alive and pre-stable with respect to E in A_K. A set of states, Q, is a pre-stabilizable set if there exists a feedback law $K(x)$ so that every $x \in Q$ is alive and pre-stable in A_K, and A is a pre-stabilizable system if X is a pre-stabilizable set. ❑

Figure 3.1 illustrates the importance of the liveness requirement in the above definition. Note that 1 is pre-stabilizable since disabling γ pre-stabilizes 1. On the other hand, disabling γ at 2 leaves no other defined event at 2. Thus, neither 2 nor 3 is pre-stabilizable.

Let Q_1 and Q_2 be two pre-stabilizable sets. Clearly, any feedback that pre-stabilizes either one of them also pre-stabilizes their intersection. Thus, pre-stabilizable sets are closed under intersections. The following result states that they are also closed under union:

Proposition 3.2 Given pre-stabilizable sets Q_1 and Q_2, $Q_1 \cup Q_2$ is also pre-stabilizable.

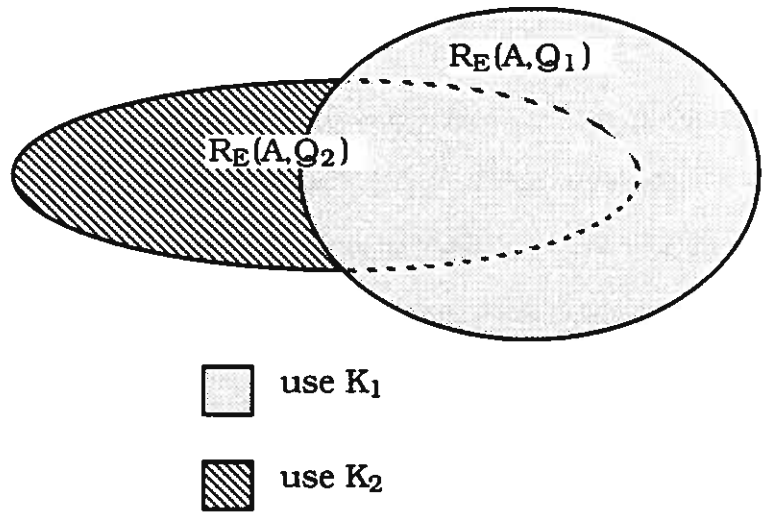

Figure 3.2: Feedback that Pre-Stabilizes the Union of Two Sets

Proof: We show this by constructing a feedback that pre-stabilizes the union. First let K_i pre-stabilize Q_i. Then, pick, say, K_1 for the reach of Q_1, and K_2 for those states in the reach of Q_2, but not in the reach of Q_1 (see Figure 3.2). More precisely, we pick a feedback F as follows:

$$F(x) = \begin{cases} K_1(x) & \text{if } x \in R_E(A_{K_1}, Q_1) \\ K_2(x) & \text{if } x \in R_E(A_{K_2}, Q_2) \cap \overline{R_E(A_{K_1}, Q_1)} \\ \text{don't care} & \text{otherwise} \end{cases}$$

Recall that $R_E(A, Q)$ is the set of states that can be reached from Q by trajectories that do not exit E once they enter it. Clearly, Q_1 is pre-stable in the closed loop system A_F. By Lemma 2.6, Q_2 is also pre-stable, since the trajectories, from a state in Q_2, either go to E or go to a state that is also in the range of Q_1, in which case, they will eventually go to E. Thus, F pre-stabilizes $Q_1 \cup Q_2$. ❑

We immediately have the following corollary:

Corollary 3.3 A maximal set that is pre-stabilizable with respect to E exists. Let $P(E)$ denote this set. ❑

Recall, from Lemma 2.5, that a necessary and sufficient condition for the pre-stability of a state x is the pre-stability of $f(x, d(x))$. A natural generalization of this condition is that the set of events defined at x can be restricted, say by $K \subset \Sigma$, so that there is at least one event defined from x, i.e., $d_K(x) = (d(x) \cap K) \cup e(x) \neq \emptyset$, and all those events take x to pre-stabilizable states, i.e., $f(x, d_K(x))$ is a pre-stabilizable set:

Lemma 3.4 $x \in X$ is pre-stabilizable with respect to E iff $x \in E$ or there exists some $K \subset \Sigma$ such that $d_K(x) = (d(x) \cap K) \cup e(x) \neq \emptyset$ and $f(x, d_K(x))$ is pre-stabilizable with respect to E.

Proof: ($\rightarrow$) Immediate from Lemma 2.5.

($\leftarrow$) Let K_1 be a feedback that pre-stabilizes $f(x, d_K(x))$. Let

$$K_2(x') = \begin{cases} K_1(x') & \text{if } x' \in R_E(A_{K_1}, f(x, d_K(x))) \\ K & \text{if } x' = x \\ \text{don't care} & \text{otherwise} \end{cases}$$

Then, also by Lemma 2.6, K_2 pre-stabilizes x. Note that we do not care about the feedback for states other than the ones we have included in the above equation for K_2, since x cannot reach those states under K_2 ❑

Now, we can construct a natural counterpart of Algorithm 2.7 using the above lemma. As in Algorthm 2.7, we start with E and then add in the states that satisfy Lemma 3.4. In particular, at each step, we include the states x for which $e(x) \neq \emptyset$ and $f(x, e(x))$ is a subset of the current set of states which are known to be pre-stabilizable, or (if $e(x) = \emptyset$), we can find an event $\sigma \in d(x)$ such that $f(x, \sigma)$ is a

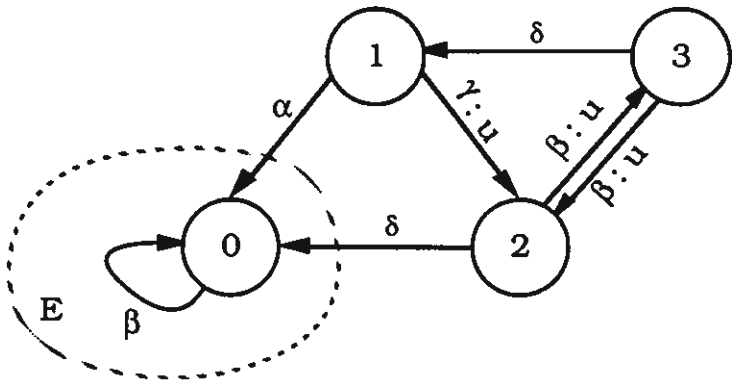

Figure 3.3: Example for the Pre-Stabilizability Algorithm

subset of the current set. For example, in Figure 3.3, We start with the state 0. At the first step, we include 1 and 2, and at the second step we include 3.

Proposition 3.5 The following algorithm computes $P(E)$ and a feedback that pre-stabilizes it. It has complexity $O(n^2)$:

Algorithm Let $X_0 = E$ and iterate:

$$P_{k+1} = \left\{ \begin{array}{ll} x| & (e(x) \neq \emptyset \text{ and } f(x, e(x)) \subset X_k) \text{ or} \\ & (e(x) = \emptyset \text{ and } \exists \sigma \in d(x) \text{ such that } f(x, \sigma) \subset X_k) \end{array} \right\}$$

$$K(x) = \left\{ \begin{array}{ll} \emptyset & \text{if } e(x) \neq \emptyset \\ \text{some } \sigma \text{ such that } f(x, \sigma) \subset X_k & \text{otherwise} \end{array} \right. , \qquad \text{for } x \in P_{k+1}$$

$$X_{k+1} = X_k \cup P_{k+1}$$

Terminate when $X_{k+1} = X_k$.

Proof: Straightforward by following the proof of Proposition 2.7 and using Lemma 3.4. ❑

In Figure 3.3, γ is disabled at 1 and β is disabled at 2 in the first step, and β is disabled at 3 in the second step.

Algorithm 3.5 leads to a feedback that disables as many events as possible. We formalize this as follows:

Definition 3.6 A pre-stabilizing feedback K is maximally restrictive, if for any pre-stabilizing feedback K' such that $d_{K'}(x) \subset d_K(x)$ for all $x \in P(E) \cap \overline{E}$, $K' = K$. ❑

We immediately have the following result:

Proposition 3.7 A pre-stabilizing feedback K is maximally restrictive iff

$$K(x) = \begin{cases} \sigma \in d(x) & \text{if } e(x) = \emptyset \\ e(x) & \text{otherwise} \end{cases},$$

for all $x \in P(E) \cap \overline{E}$.

Proof: Straightforward. ❑

Thus, Algorithm 3.5 leads to a maximally restrictive feedback.

The feedback of Algorithm 3.5, also minimizes the maximum number of transitions it takes to go from a state to E. Clearly, it also minimizes the radius. In Figure 3.3, it takes a single transition to go from 1 or 2 to 0, and two transitions to go from 3 to 0. To formalize this, let us assume, without loss of generality, that A is stabilizable, and let $r(A, x)$ denote the length of the longest path from x to a state in E, where $r(A, x) = 0$ for all $x \in E$:

Definition 3.8 Given a stabilizable system A, a pre-stabilizing state feedback K is path minimizing if for any pre-stabilizing feedback K' such that $r(A_{K'}, x) \leq r(A_K, x)$ for all x, $r(A_{K'}, x) = r(A_K, x)$ for all x. ❑

For any two path-minimizing feedbacks, the longest path lengths are equal at each state:

Proposition 3.9 For any two path minimizing feedbacks K_1, and K_2, $r(A_{K_1}, x) = r(A_{K_2}, x)$ for all x.

Proof: We prove this by assuming the contrary and contructing a feedback, F which leads to shorter path lengths than those of K_1 and K_2.

If we assume the contrary, then there exists $x, y \in X$ such that $r(A_{K_1}, x) < r(A_{K_2}, x)$ and $r(A_{K_1}, y) > r(A_{K_2}, y)$. To construct F, we want to use K_1 for the reach of x and K_2 for the reach of y. However, we need to make sure that we do not introduce any loops by doing this. Thus, we want to show that if $x \in R_E(A_{K_2}, y)$ then $y \notin R_E(A_{K_1}, x)$, and vice versa.

Assume the contrary to this, i.e., assume that $x \in R_E(A_{K_2}, y)$, and $y \in R_E(A_{K_1}, x)$. If $x \in R_E(A_{K_2}, y)$, then $r(A_{K_2}, y) > r(A_{K_2}, x)$. On the other hand, if $y \in R_E(A_{K_1}, x)$, then $r(A_{K_1}, x) > r(A_{K_1}, y)$. Since also $r(A_{K_1}, x) < r(A_{K_2}, x)$ and $r(A_{K_1}, y) > r(A_{K_2}, y)$, $r(A_{K_2}, x) > r(A_{K_2}, y)$. Thus, we have a contradiction.

Therefore, the following choice of F is pre-stabilizing, if we assume, without loss of generality, that $x \in R_E(A_{K_2}, y)$:

$$F(x') = \begin{cases} K_1(x') & \text{if } x' \in R_E(A_{K_x}, x) \\ K_2(x') & \text{otherwise} \end{cases}$$

F also leads to shorter path lengths than those of K_1 and K_2, and thus, neither K_1 nor K_2 is path-minimizing. Therefore, $r(A_{K_1}, x) = r(A_{K_2}, x)$ for all x. ❑

Let $\overline{r}(x)$ denote this minimal path length for x. Finally, we show that Algorithm 3.5 leads to a path-minimizing feedback, and it also constructs $\overline{r}(x)$:

Proposition 3.10 The feedback of Algorithm 3.5, K, is path-minimizing. Furthermore, $x \in P_k$ at step k of the algorithm, if and only if $k = \overline{r}(x)$, and thus, the algorithm also constructs $\overline{r}(x)$.

Proof: Suppose that K is path minimizing for X_k at step k of Algorithm 3.5. Then, for some $x \in P_{k+1}$, $f(x, d_K(x)) \cap P_k \neq \emptyset$ since otherwise $x \in P_i$ for some $i < k+1$. Therefore, K is path minimizing for X_{k+1} and by induction, K is path-minimizing for $P(E)$. The proof of the second statement is straightforward. ❑

If all the trajectories in the desired behavior consist of states in E, then a maximally restrictive feedback is desirable for stabilization since it does not restict the desired behavior, and in addition, it ensures returning to E in a minimum number of transitions. However, if the desired behavior involves states outside of E (for example, if E is simply a set of desired initial states), then one would prefer a less restrictive feedback so that all stable trajectories of the desired behavior are enabled. In what follows, we present an algorithm to construct a feedback that disables as few events as possible:

Definition 3.11 A pre-stabilizing feedback K is minimally restrictive, if for any pre-stabilizing feedback K' such that $d_{K'}(x) \supset d_K(x)$ for all $x \in P(E) \cap \overline{E}$, $K' = K$. ❑

Our algorithm is based on the following lemma, which states that a feedback, K, is minimally restrictive if and only if enabling any event at any state, which is otherwise disabled, makes that state unstable, i.e., creates a cycle that does not go through E. That is, K is minimally restrictive if for any state x, enabling any σ that had been disabled by K can move the system to a new state from which it can return to x without going through E. For example, in Figure 3.3, consider the feedback which disables γ at 1 and β at 3. Enabling either of these events makes the corresponding state unstable.

Lemma 3.12 A pre-stabilizing state feedback K is minimally restrictive iff for all $x \in P(E)$ and $\sigma \in d(x) \cap \overline{K(x)} \cap \overline{e(x)}$, $x \in R_E(A_K, f(x, \sigma))$.

Proof: ($\rightarrow$) Straightforward by assuming the contrary.

($\leftarrow$) Since we cannot enable any event without making some state unstable, K is certainly a minimally restrictive feedback. ❑

In order to compute a minimally restrictive feedback, we start with a maximally restrictive feedback, and add events, that are otherwise disabled, until the condition of the the above lemma is satisfied. Our algorithm visits all states in $P(E) \cap \overline{E}$ and for each state x, it includes all events $\sigma \in d(x) \cap \overline{K(x)} \cap \overline{e(x)}$, such that $x \notin R_E(A_K, f(x,\sigma))$, in $K(x)$. Since K is possibly modified after visiting a state, when the next state is visited, the new feedback should be used in computing $R_E(A_K, f(x,\sigma))$. For example, in Figure 3.3, if we start with state 3, we get the minimally restrictive feedback which disables γ at 1 and β at 3. Depending on the order the states are visited, different minimally restrictive feedbacks may be generated. In Figure 3.3, if we start with state 1 or 2, we get the minimally restrictive feedback which disables β at 2.

On the other hand, we do not need to compute $R_E(A_K, f(x,\sigma))$ for each σ and x. Instead, we can compute, for each x, the set of states that <u>can</u> reach x and check to see if any element of $f(x,\sigma)$ is in this set:

Lemma 3.13 Given x and $\sigma \in d(x) \cap \overline{K(x)} \cap \overline{e(x)}$, $x \in R_E(A_K, f(x,\sigma))$ iff $f(x,\sigma) \cap R_E(A_K^{-1}, x) \neq \emptyset$.

Proof: Straightforward. ❑

We then have the following algorithm:

Proposition 3.14 The following algorithm computes a minimally restrictive feedback. It has complexity $O(n^2)$:

Algorithm For all $x \in P(E)$ do:

$$\begin{aligned} S &= K(x) \cup \{\sigma \in d(x) \cap \overline{K(x)} \cap \overline{e(x)} | f(x,\sigma) \cap R_E(A_K^{-1}, x) = \emptyset\} \\ K(x) &= S \end{aligned}$$

Proof: The proof follows from Lemma 3.12. The complexity is $O(n^2)$ since all states are visited and the reach operation has complexity $O(n)$. ❑

3.2 Stabilizability and (f, u)-Invariance

Stabilizability, like pre-stabilizability, is defined as a natural extension of stability. A state x is stabilizable if we can find a state feedback such that x is stable in the closed loop system:

Definition 3.15 Given a live system A and some $E \subset X$, $x \in X$ is <u>stabilizable</u> with respect to E if there exits a state feedback K such that x is alive and stable with respect to E in A_K. A set of states, Q, is a <u>stabilizable set</u> if there exists a feedback law $K(x)$ so that every $x \in Q$ is alive and stable in A_K, and A is a <u>stabilizable system</u> if X is a pre-stabilizable set. ❑

Let Q_1 and Q_2 be two stabilizable sets. Clearly, any feedback that stabilizes either one of them also stabilizes their intersection. Thus, stabilizable sets are closed under intersections. The following result states that they are also closed under union:

Proposition 3.16 Given stabilizable sets Q_1 and Q_2, $Q_1 \cup Q_2$ is also stabilizable.

Proof: We show this by constructing a feedback that stabilizes the union. First let K_i stabilize Q_i. Then, pick, say, K_1 for the reach of Q_1, and K_2 for those states in

the reach of Q_2, but not in the reach of Q_1. More precisely, we pick a feedback F as follows:

$$F(x) = \begin{cases} K_1(x) & \text{if } x \in R(A_{K_1}, Q_1) \\ K_2(x) & \text{if } x \in R(A_{K_2}, Q_2) \cap \overline{R(A_{K_1}, Q_1)} \\ \text{don't care} & \text{otherwise} \end{cases}$$

F clearly stabilizes $Q_1 \cup Q_2$. ❑

We immediately have the following corollary:

Corollary 3.17 A maximal set that is stabilizable with respect to E exists. Let $S(E)$ denote this set. ❑

To achieve stabilizability, we try to make $P(E)$ f-invariant by choice of feedback. Thus, we first need to define the following counterpart of the notion of (A, B)-invariance:

Definition 3.18 A subset Q of X is <u>(f, u)-invariant</u> if there exists a state feedback K such that Q is f-invariant in A_K. ❑

Let $A_\emptyset$ denote A with all controllable events disabled, i.e., $A_\emptyset = (X, f, e, e)$. Note that if some Q is f-invariant in $A_\emptyset$ then it is also (f, u)-invariant in A. The following result formalizes this and it also establishes a connection with the well-known result in [24] that a subspace $\mathcal{V}$ is (A, B)-invariant iff $A\mathcal{V} \subset \mathcal{V} + \mathcal{B}$, where $\mathcal{B}$ is the range of B (compare to item 2 below):

Proposition 3.19 The following statements are equivalent:

1. Q is (f, u)-invariant in A.
2. $\forall x \in Q,\ f(x, d(x)) \subset Q \cup \overline{f(x, e(x))}$.

3. Q is f-invariant in $A_\emptyset$.

Proof: ($1 \to 2$) Suppose that there exists $x \in Q$, $\sigma \in d(x)$, $y \in f(x,\sigma)$ such that $y \notin Q$ and $y \notin \overline{f(x,e(x))}$. But, then $y \in f(x,e(x))$ and we have a contradiction since the transition to y is undesired and cannot be disabled.

($2 \to 3$) Assume the contrary. By Proposition 2.13, there exists $x \in Q$, $\sigma \in e(x)$ such that $f(x,\sigma) \not\subset Q$. By (2), $f(x,\sigma) \subset \overline{f(x,e(x))}$ and we have a contradiction.

($3 \to 1$) Simply use the feedback $K(x) = e(x)$. ❑

The class of (f,u)-invariant sets is closed under arbitrary unions and intersections. Thus, a unique maximal (f,u)-invariant subset of Q exits, and let $T(Q)$ denote this set. Clearly, $T(Q) = \overline{R(A_\emptyset^{-1}, \overline{Q})}$. This notion of (f,u)-invariance, however, is not sufficient for stabilizablity since we also need to keep $S(E)$ alive. So, we define the following notion which requires that we can find a state feedback such that Q is both alive and f-invariant in the closed loop system:

Definition 3.20 A subset Q of X is a <u>sustainably (f,u)-invariant</u> set if there exists a state feedback K such that Q is alive and f-invariant in A_K. ❑

The class of sustainably (f,u)-invariant sets is closed under arbitrary unions but not intersections. In Figure 3.4, Q_1 (respectively Q_2) can be made (f,u)-invariant and alive by disabling δ (respectively α) at state 1. On the other hand, $Q_1 \cup Q_2$ is clearly sustainably (f,u)-invariant, but $Q_1 \cap Q_2$ is not, since if both α and δ are disabled, the state 1 is no longer alive. Let $I(Q)$ denote the maximal sustainable (f,u)-invariant subset of Q.

The characterization of sustainable (f,u)-invariance requires a slightly more careful look at what it means if an (f,u)-invariant set is <u>not</u> sustainable. Specifically,

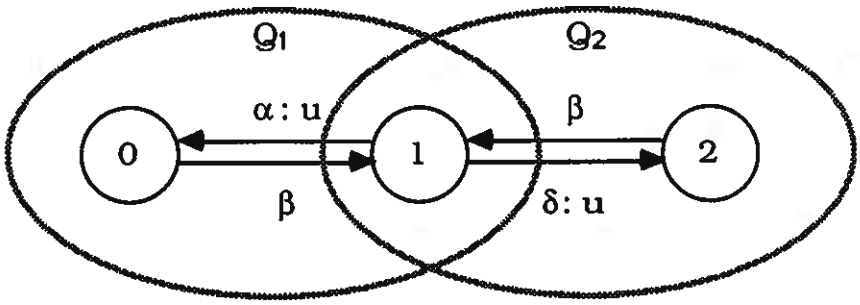

Figure 3.4: Example for the Sustainable (f, u)-Invariance of Unions and Intersections

for any Q and all states x in $T(Q)$, we know that all events that may take x outside of $T(Q)$ are controllable, and can therefore be disabled to achieve the desired invariance. However, x may have no other events defined, and thus, making $T(Q)$ f-invariant will disable all events from x. Then, $T(Q)$ is not sustainable. Our algorithm for computing $I(Q)$ is based on first computing $T(Q)$ and then throwing away all states that are no longer alive. We then apply the same procedure to this new set. This iteration continues until no states are discarded on a step. The following result states that we then have a sustainable (f, u)-invariant set:

Proposition 3.21 Given $Q \subset X$, let

$$Q' = T(\{x \in Q | \text{ there exists some } \sigma \in d(x) \text{ such that } f(x, \sigma) \subset Q\})$$

Then, Q is sustainably (f, u)-invariant iff $Q' = Q$.

Proof: $(\rightarrow)$ Obvious

$(\leftarrow)$ If $Q' = Q$ then $Q = T(Q)$ and for all $x \in Q$ there exists some $\sigma \in d(x)$ such that $f(x, \sigma) \subset Q$. Therefore, Q is alive and f-invariant in A_K with $K(x) = \{\sigma | f(x, \sigma) \subset Q\}$. ❑

This result implies that in order to find $I(Q)$, we can apply the operation T iteratively by throwing away the states that are no longer alive after the last step:

Proposition 3.22 The following algorithm computes $I(Q)$ and it has complexity $O(n^2)$:
Algorithm Let $X_0 = Q$. Iterate:

$$X_{k+1} = T(\{x \in X_k | \text{ there exists } \sigma \in d(x) \text{ such that } f(x,\sigma) \subset X_k\})$$

Terminate when $X_{k+1} = X_k$.

Proof: Clearly, this algorithm terminates in a finite number of steps, say r steps. By Lemma 3.21, $X_r \subset I(Q)$. On the other hand, $I(Q) \subset X_0$. Suppose that $I(Q) \subset X_k$ for some k. Then, $X_{k+1} \supset T(I(Q)) = I(Q)$. Thus, $I(Q) \subset X_k$ for all k, implying that $I(Q) = X_r$. To justify the computational complexity, note that we visit each state at most once at each iteration, and there can be at most n iterations. ❑

Now, we proceed with deriving an algorithm for the maximal stabilizable set, $S(E)$. We begin by computing $P(E)$, the maximal pre-stabilizable set with respect to E. If $P(E)$ were sustainably (f,u)-invariant, we could be done, with $S(E) = P(E)$. More generally, however, there may be some states in $P(E)$ for which it is impossible to find a feedback which keeps trajectories within $P(E)$. Furthermore, since all elements of $P(E)$ are pre-stabilizable, some of these troublesome states <u>must</u> be in E. Thus, what we must do is to compute $I(P(E))$ and then discard elements of E <u>not</u> in $I(P(E))$, reducing E to a new set E'. However, there may now be states which were pre-stabilizable with respect to E but <u>not</u> with respect to E', and we therefore must repeat the process. For example, in Figure 3.5, $E = \{0,3\}$, $P(E) = \{0,1,2,3\}$ and $I(P(E)) = \{1,2,3\}$, so that $E' = \{3\}$. Now, $\{1\}$ is <u>not</u> pre-stabilizable with respect to E' and must be discarded. The iteration in this case would produce

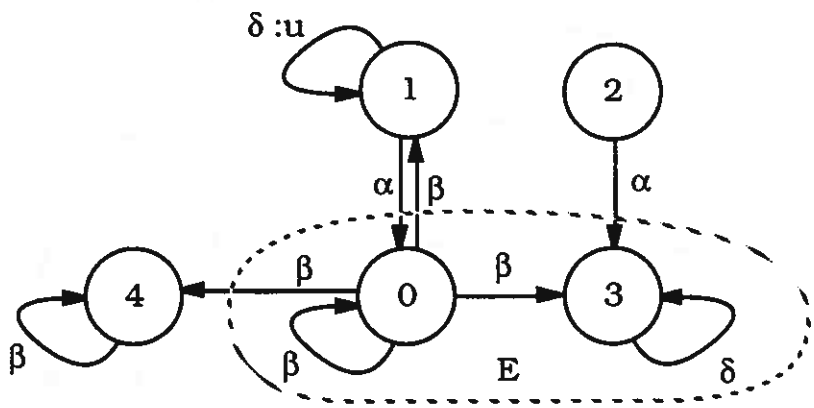

Figure 3.5: Example for the Stabilizability Algorithm

$P(E') = \{2,3\} = I(P(E'))$. It is not difficult to check that $\{2,3\} = S(E)$ as well, and indeed the following result states this more generally and provides the basis for our algorithm:

Lemma 3.23 Given $E, Q \subset X$, let

$$Q' = I(P(E \cap Q))$$

If $Q' = Q$ then Q is stabilizable with respect to E.

Proof: We first show that $Q = P(E \cap Q)$. If $Q' = Q$, then, clearly, $Q \subset P(E \cap Q)$. In order to show that $Q \supset P(E \cap Q)$, suppose that there exists some $x \in P(E \cap Q)$ such that $x \notin Q$ (see Figure 3.6). Then, there exists some feedback such that all paths from x go to $E \cap Q$ with trajectories that only lie in $P(E \cap Q)$. Since by assumption $Q = I(P(E \cap Q))$ but $x \notin Q$, x can reach some $y \in E \cap Q$ and $\sigma \in e(y)$ such that $f(y, \sigma) \not\subset P(E \cap Q)$. But then, Q cannot be sustainably (f, u)-invariant, which is a contradiction. Therefore, $Q = P(E \cap Q)$, and this immediately implies that $Q = I(Q)$ as well.

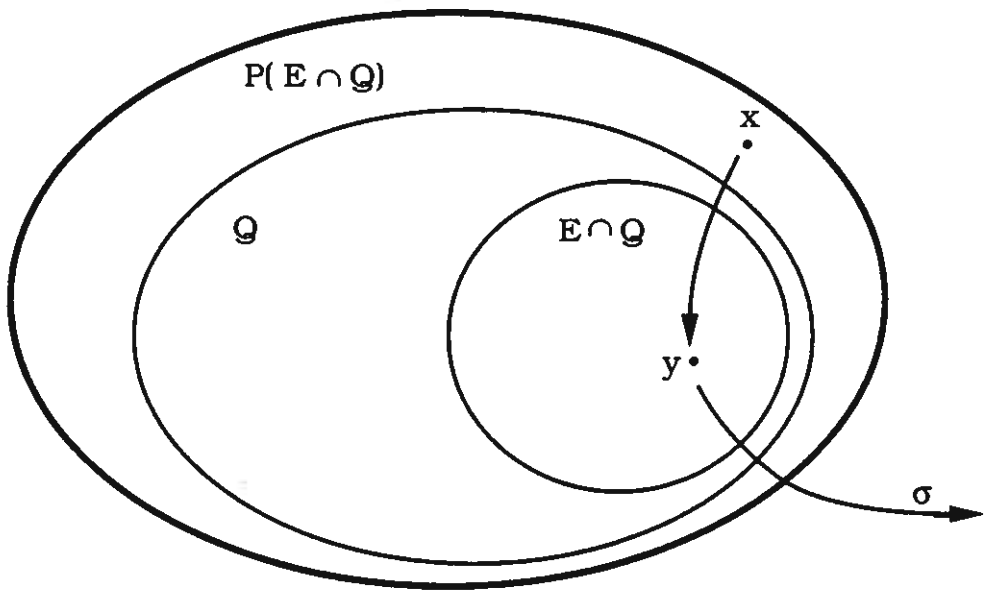

Figure 3.6: Illustration for the Proof of Lemma 3.23

In order to show that Q is stabilizable, let K_P denote a feedback that pre-stabilizes $P(E \cap Q)$ and let

$$K(x) = \begin{cases} K_P(x) & \text{if } x \in \overline{E} \cap Q \\ \{\sigma | f(x,\sigma) \subset Q\} & \text{if } x \in E \cap Q \end{cases}$$

Q is clearly alive in A_K. Q is also pre-stable in A_K since K_P insures that all paths from states in $Q \cap \overline{E}$ go to $Q \cap E$. Q is then stable in A_K. Therefore, Q is stabilizable, and K is a stabilizing feedback. ❑

This result leads to the following algorithm:

Proposition 3.24 The following algorithm computes $S(E)$ and a feedback that stabilizes it. It has complexity $O(n^3)$.

Algorithm Let $X_0 = X$ and iterate:

$$X_{k+1} = I(P(E \cap X_k))$$

Terminate when $X_{k+1} = X_k$.

Proof: To show that $X_{k+1} \subset X_k$ for all k, first note that $X_1 \subset X_0$. Assume that $X_{k+1} \subset X_k$ for some k. Let P_k denote $P(E \cap X_k)$. Note that $P_{k+1} \subset P_k$ since $E \cap X_{k+1} \subset E \cap X_k$. Then, $I(P_{k+1}) \subset I(P_k)$ and thus $X_{k+2} \subset X_{k+1}$. Therefore, this algorithm terminates in a finite number of steps, say r steps. By Lemma 3.23, $X_r \subset S(E)$. On the other hand, $S(E) \subset X_0$. Assume that $S(E) \subset X_k$ for some k. Then, clearly, $S(E) \subset P(E \cap X_k)$. Since also $I(S(E)) = S(E)$, $S(E) \subset X_{k+1}$, so that $S(E) = X_r$. The feedback which achieves pre-stability at step r, say F, also achieves stability for $S(E)$, since states in $\overline{S(E)}$ are not pre-stable with respect to $E \cap S(E)$ and thus F cannot be enabling any event that takes a state in $S(E)$ outside of $S(E)$. Also, for states in $E \cap S(E)$, all events that take those states outside of $S(E)$ should be disabled. In short, the stabilizing feedback K is

$$K(x) = \begin{cases} F(x) & \text{if } x \in S(E) \cap \overline{E} \\ \{\sigma \in d(x) | f(x,\sigma) \subset S(E)\} & \text{if } x \in S(E) \cap E \\ \text{don't care} & \text{otherwise} \end{cases}$$

To justify computational complexity, recall that the computation of $P(E)$ is $O(n^2)$, and note that the above algorithm terminates in at most n steps. ❑

Note that a stabilizing feedback has two components: one component for pre-stability and another for invariance. Clearly, there is no flexibility in choosing the feedback to achieve invariance (events that do <u>not</u> take the state out of P_{k+1} should <u>not</u> be disabled of course). For pre-stability either a minimally restrictive or a maximally

restrictive feedback can be chosen. Then, the corresponding stabilizing feedback is, accordingly, minimally restrictive or maximally restrictive.

4 Conclusions

In this paper, we have introduced notions of stability and stabilizability for discrete-event systems described by finite-state automata, and we have developed polynomial algorithms to test for stability and stabilizability and to construct maximal stable and stabilizable sets, and for the latter, a feedback control law that makes a stabilizable set or system stable. Our work has drawn on a blend of concepts from computer science and from dynamic systems and control. In particular, the notion of pre-stability used here is well-known in the computer science literature, while the concepts of state feedback, f-invariance, and (f, u)-invariance that are of critical importance for our study of stability and stabilizability, are control concepts in systems and control.

The stability concepts that we introduced here can be thought of as notions of error recovery or resiliency in that the system always returns to "good" states. From the control perspective, one can also formulate several related concepts and problems. For example, in many applications, one may not have (or want) access to full state or event information but may still wish to stabilize the system. This leads directly to questions of state reconstruction, observability, and output feedback. Also, motivated by problems such as schedule-following in a flexible manufacturing system, one can formulate regulator or tracking problems for DEDS in which a feedback system is sought so that the DEDS produces a particular desired sequence of output events. Analysis addressing these and related problems will be the subjects of subsequent papers.

Acknowledgements

The authors wish to acknowledge Professors George C. Verghese and Gilead Tad-

more for their insightful comments in a preliminary phase of this work. Authors Professor Willsky, and Mr. Özveren also wish to acknowledge Institut de Recherche en Informatique et Systèmes Aléatoires, Rennes, France, for their support during our stay, and Professors Jean-Claude Raoult, Albert Benveniste, Mr. Paul Le Guernic, and Mr. Bernard Le Goff for their comments on this work.

References

[1] M. Ben-Ari, *Principles of Concurrent Programming*, Prentice Hall, 1982.

[2] G.V. Bochmann, *Distributed System Design*, Springer Verlag, 1983.

[3] R. Cieslak, C. Desclaux, A. Fawaz, P. Variaya, "Modeling and Control of Discrete Event Systems", *Proceedings of CDC*, December 1986.

[4] E.M. Clarke, O. Grumberg, "Research on Automatic Verification of Finite-State Concurrent Systems", *CMU-CS-87-105*, Jan. 1987.

[5] E.W. Dijkstra, "Solution of a Problem in Concurrent Programming Cotrol", *CACM*, Vol. 5, No. 9, Sept. 1965.

[6] E.W. Dijkstra, "Self-Stabilizing Systems om Spite of Distributed Control", *CACM*, Vol. 17, No. 11, Nov. 1974, pp. 643-644.

[7] G.A. Frank, D.L. Franke, W.F. Ingogly, "An Architecture Design and Assessment System", *VLSI Design*, August 1985.

[8] C.H. Golazewski, P.J. Ramadge, "Mutual Exclusion Problems for Discrete Event Systems with Shared Events", *Conference on Decision and Control*, Dec. 1988.

[9] W.B. Gevarter, "Expert Systems: Limited but Powerful", *IEEE Spectrum*, August 1983.

[10] F. Lin, W.M. Wonham, "Decentralized Supervisory Control of Discrete Event Systems", *Systems Control Group Report 8612, University of Toronto*, July 1986.

[11] O.Z. Maimon, G. Tadmor, "Efficient Supervisors in Discrete Event Systems", *Proceedings of 1986 Int. Conf. Systems, Man and Cybernetics.*

[12] M.E. Merchant, "Production: A Dynamic Challenge", *IEEE Spectrum*, May 1983.

[13] J.S. Ostroff, W.M. Wonham, "A Temporal Logic Approach to Real Time Control", *Proceedings of CDC*, December 1985.

[14] C.M. Özveren, "Analysis and Control of Discrete Event Dynamic Systems: A State Space Approach", *PhD Thesis*, in preperation.

[15] P.J. Ramadge, W.M. Wonham, "Supervisory Control of a Class of Discrete Event Processes", *SIAM J. of Cont. and Opt.*, January 1987.

[16] P.J. Ramadge, W.M. Wonham, "Modular Feedback Logic for Discrete Event Systems", *SIAM J. of Cont. and Opt.*, September 1987.

[17] J. Sifakis, "Deadlocks and Livelocks in Transition Systems", R.R. No:185, January 1980.

[18] G. Tadmor, O.Z. Maimon, "Control of Large Discrete Event Systems: Constructive Algorithms", *MIT LIDS Publication*, LIDS-P-1627, December 1986.

[19] J.G. Thistle, W.M. Wonham, "Control Problems in a Temporal Logic Framework", *Int. J. Control 44 (4)*, 1986.

[20] W. Thomas, "Automata on Infinite Objects", *Lehrstuhl für Informatik II*, RWTH Aachen, D-5100 Aachen, April 1988.

[21] L. Tobias, J.L. Scoggins, "Time-Based Air-Traffic Management Using Expert Systems", *IEEE Control Systems Magazine*, April 1987.

[22] A.F. Vaz, W.M. Wonham, "On Supervisor Reduction in Discrete Event Systems", *International Journal of Control*, 1986.

[23] W.M. Wonham, P.J. Ramadge, "On the Supremal Controllable Sublanguage of a Given Language", *SIAM J. of Cont. and Opt.*, May 1987.

[24] W.M. Wonham, *Linear Multivariable Control: A Geometric Approach*, Springer-Verlag Inc., New York, 1985.

70

JACM91