lookup procedures, usually augmented with special devices to handle phonological changes that operate across word boundaries. Another approach to solving this computational problem would be to use the reversed cascade of rules during recognition, but to somehow make the filtering information of particular rules available earlier in the process. However, no general and effective techniques have been proposed for doing this.

The more radical approach that we explore in this paper is to eliminate the cascade altogether, representing the information in the grammar as a whole in a single more unified device, namely, a finite-state transducer. This device is constructed in two phases. The first is to create for each rule in the grammar a transducer that exactly models its behavior. The second is to compose these individual rule transducers into a single machine that models the grammar as a whole.

Johnson (1972) was the first to notice that the noncyclic components of standard phonological formalisms, and particularly the formalism of *The Sound Pattern of English* (Chomsky and Halle 1968), were equivalent in power to finite-state devices despite a superficial resemblance to general rewriting systems. Phonologists in the SPE tradition, as well as the structuralists that preceded them, had apparently honored an injunction against rules that rewrite their own output but still allowed the output of a rule to serve as context for a reapplication of that same rule. Johnson realized that this was the key to limiting the power of systems of phonological rules. He also realized that basic rewriting rules were subject to many alternative modes of application offering different expressive possibilities to the linguist. He showed that phonological grammars under most reasonable modes of application remain within the finite-state paradigm.

We observed independently the basic connections between rewriting-rule grammars and finite-state transducers in the late 1970s and reported them at the 1981 meeting of the Linguistic Society of America (Kaplan and Kay 1981). The mathematical analysis in terms of regular relations emerged somewhat later. Aspects of that analysis and its extension to two-level systems were presented at conferences by Kaplan (1984, 1985, 1988), in courses at the 1987 and 1991 Linguistics Institutes, and at colloquia at Stanford University, Brown University, the University of Rochester, and the University of Helsinki.

Our approach differs from Johnson's in two important ways. First, we abstract away from the many details of both notation and machine description that are crucial to Johnson's method of argumentation. Instead, we rely strongly on closure properties in the underlying algebra of regular relations to establish the major result that phonological rewriting systems denote such sets of string-pairs. We then use the correspondence between regular relations and finite-state transducers to develop a constructive relationship between rewriting rules and transducers. This is accomplished by means of a small set of simple operations, each of which implements a simple mathematical fact about regular languages, regular relations, or both. Second, our more abstract perspective provides a general framework within which to treat other phonological formalisms, existing or yet to be devised. For example, two-level morphology (Koskenniemi 1983), which evolved from our early considerations of rewriting rules, relies for its analysis and implementation on the same algebraic techniques. We are also encouraged by initial successes in adapting these techniques to the autosegmental formalism described by Kay (1987).

2. Rewriting Rules and Transducers

Supposing for the moment that Rule 2 ($N \rightarrow n$) is optional, Figure 1 shows the transition diagram of a finite-state transducer that models it. A finite-state transducer has