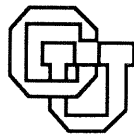An Evolving Simulation of
the Gries/Dijkstra Design Process

Robert B. Terwilliger

CU-CS-632-92          December 1992

University of Colorado at Boulder
DEPARTMENT OF COMPUTER SCIENCE

# An Evolving Simulation of
# the Gries/Dijkstra Design Process

Robert B. Terwilliger

Department of Computer Science
University of Colorado
Boulder, CO 80309-0430
*email*: 'terwilli@cs.colorado.edu'

## ABSTRACT

We are applying a three part, investigative approach to the formal design method described by Gries in his book *The Science of Programming*. First, we walkthrough the process on a number of problems; second, we construct a simulation program which duplicates the designs produced by the walkthroughs; and third, we produce a process program that supports human application of the method. In this paper, we describe the evolution of our simulation program through two iterations of this three step procedure. In our system, design knowledge is stored as cliches: complex structures representing commonly occurring situations. Development consists of applying cliches to specifications until a complete design is produced. We have construted two distinct cliche libraries which differ significantly in their description of the design process.

## 1. Introduction

Some have suggested that *formal methods* can enhance both the software specification and design processes [2-5, 8, 14, 19, 21, 22, 25, 30, 44]. For example, Dijkstra and Gries [9, 10, 16, 17] have developed a process which takes a pre- and post-condition specification written in first-order predicate logic and incrementally transforms it into a verified design written using guarded commands. Although this technique has been used and taught for over a decade, at present is it difficult to either correctly apply it, determine if it has been applied properly, or evaluate it for effectiveness.

One approach to gathering information about development processes is the use of *walkthroughs* and *inspections* [11, 13, 42, 45]. In these situations, a software item or the method used to produce it is presented to a group who evaluate it according to an appropriate set of criteria. Even more information can be gathered through *process programming* [18, 26, 33, 34]: describing software processes using programming language constructs and notations. To a certain extent, this effort has similarities to previous work on *knowledge-based software engineering* [1, 12, 15, 20, 23, 24, 27, 28, 31]. An important contribution of this community is the notion of *cliche* (or *plan* or *schema*): a complex knowledge structure representing a commonly occurring situation.

We are applying a three part, investigative approach to the formal design method described by Gries in his book *The Science of Programming*. First, we walkthrough, in other words hand simulate, the process on a number of problems. This produces an increased understanding of the method as well as a suite of example designs. Second, we produce a program that simulates the design process discovered during the walkthroughs; ideally, it should be able to recreate the suite of designs previously produced. Third, we produce a process program that supports human application of the method.

We have currently completed two iterations of our three part approach on the Gries/Dijkstra process [37-41]; our understanding of the method has evolved significantly during the time we have spent studying it. After performing the initial walkthroughs, we viewed the design method as a linear sequence of small, relatively independent steps; however, our attempts to automate the process based on this model met with no success. Instead, both our simulation and process programs are based on a library of relatively large cliches representing solutions to common design problems.

Our simulation of the Gries/Dijkstra process has two levels. At the lower level, a design derivation sub-process transforms formal specifications into correct designs using a library of cliches. On the upper level, a cliche derivation sub-process uses strategies and proof rules to construct and verify cliches. The advantage of our two level architecture is that proofs are performed mostly at "compile" rather than "run" time. Cliche construction and verification is quite difficult, but is done only once for each cliche and performed by a human. On the other hand, cliche application is reasonably easy and is performed repeatedly by the machine.

We have produced two distinct cliche libraries which differ significantly in their description of the design process. Our first iteration cliches combine a number of what were considered distinct steps during the initial walkthroughs into a single unit. Our second iteration cliches combine several steps from the initial

walkthroughs, but not nearly as many as the first iteration cliches do. They also better reflect the process described by Gries [17]. We believe our second iteration cliches are a significant step towards an accurate and detailed description of the method.

In the remainder of this paper, we describe the evolution of our simulation program in more detail. In section two we describe the Gries/Dijkstra design process as we understood it after our initial walkthroughs, and in section three we discuss the architecture of our simulation system. In section four, we give an example of our first iteration cliches and their use in design derivation, and in section five we do the same for the second iteration. Finally, in section six, we summarize and draw some conclusions from our experience.

## 2. Gries/Dijkstra Design

Figure 1 shows a pictorial representation of the method described in [17] as we understood it after completing our initial walkthroughs. The design derivation process uses stepwise refinement to transform pre- and post-condition specifications written in first-order predicate logic into verified programs written using guarded commands. At each step, *strategies* determine how the current partial program is to be elaborated, and *proof rules* are used to verify the correctness of the transformation.

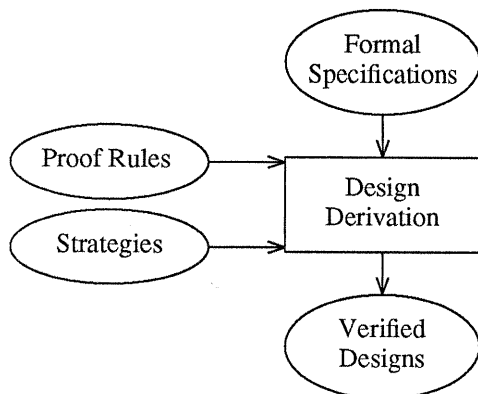For example, loops are specified using a predicate called the *invariant*, which must be true both before and after each iteration of the loop, and an integer function called the *bound*, which is an upper limit on the number of iterations remaining. The proof rule for loops has five conditions for correctness [17]:

1) the invariant must be initialized correctly
2) execution of the loop body must maintain the invariant
3) termination of the loop with the invariant true must guarantee the post-condition
4) the bound function must be greater than zero while the loop is running
5) execution of the loop body must decrease the bound.

A loop development strategy closely parallels the proof rule; each step in the strategy suggests an action, whose result can be verified using one of the items from the proof rule.

To develop a loop given a pre- and post-condition:

1) weaken the post-condition to obtain an invariant and develop a statement to initialize it
   - check item one of proof rule

2) develop a loop guard
   - check item three of proof rule

3) develop a bound function
   - check item four of proof rule

4) develop a statement that decreases the bound
   - check item five of proof rule

5) develop a statement that restores the invariant
   - check item two of proof rule.

The first step is to develop an invariant by *weakening* the post-condition; in other words, the invariant is an easier to satisfy version of the desired result. There are at least three ways to weaken the post-condition: delete a conjunct, replace a constant by a variable, and enlarge the range of a variable. An initialization for the loop must also be developed and then checked using item one of the proof rule. A good invariant has a simple initialization; if none can be found then the invariant is discarded and another weakening method is tried.

The second step is to develop a loop guard; in general, the guard will follow directly from the weakening method. For example, if the invariant is created by deleting a conjunct from the post-condition, then the guard will be the negation of the deleted conjunct. The guard's correctness is verified using item three of the proof rule.

The third step is to develop a bound function; this is done by discovering a property that should be decreased by each iteration of the loop body and then formalizing it. Item three of the proof rule is used to verify that the invariant, guard and bound are



Figure 1.  Gries/Dijkstra Design Process

consistent.

The fourth step is to develop a statement that decreases the bound; this will form the first part of the loop body. Item five of the proof rule is used to check that the bound is diminished. The final step is to develop a statement that restores the invariant after the statement which decreases the bound has invalidated it. Item two of the rule is used to check that the loop body maintains the invariant.

The above strategy is a step by step procedure for transforming a pre- and post-condition into a provably correct loop. Each step in the strategy is checked using the appropriate item from the proof rule; the loop and its proof of correctness are developed simultaneously. We can further clarify this process in the context of an example.

Kemmerer's Library problem has received considerable attention in the software engineering literature [43]. The problem is concerned with a small library database that provides both query and update transactions to library staff and users. The architectural design for our solution [37] consists of a single module that encapsulates the database and provides an entry routine for each transaction. The state of the module is modeled abstractly using high-level data types, and the entry routines are specified using pre- and post-conditions.

For example, consider the "isout" function, which returns true if a particular book is currently checked out.

```
function isout(b:book) : boolean ;
    pre   true ;
    post  isout = (∃u∈users:chkdout(u,b));
```

The variable "users" is a set that holds a record for each staff member or patron of the library. The boolean function "chkdout" takes a user and book as arguments and returns true only if the book is currently checked out to the user. The "isout" function takes a single book as an argument. The pre-condition is simply "true", while the post-condition states that the return value is true if some user has the book in question checked out.

Using the Gries/Dijkstra process, design might proceed as follows. First, we replace the constant "users" in the post-condition with the variable (expression) "users-usrs" to obtain the following invariant.

```
usrs⊆users ∧
isout = (∃u∈users-usrs:chkdout(u,b))
```

The variable "usrs" holds the users still to be examined. At any point during the loops execution,

"usrs" is a subset of "users" and "isout" is true if one of the users already examined has the book in question checked out. The invariant is initialized with the assignment "isout,usrs:=false,users", and this satisfies item one of the proof rule.

We now develop a guard for the loop. Item three of the proof rule tells us that the negation of the guard and the invariant together must imply the post-condition. Since we created the invariant from the post-condition by replacing a constant with a variable, the guard is just that the variable does not equal the constant. In our example, the loop should stop when "users-usrs" is equal to "users"; therefore, the guard is "usrs≠{}", and it satisfies item three.

We now develop a bound function. In our example, each iteration should decrease the number of elements in "usrs"; we formalize this as "|usrs|" and check that it satisfies item four of the proof rule. We now have the following.

```
{Q: true}
var usrs : set(user) ;
isout,usrs:=false,users ;
{inv P:usrs⊆users ∧
       isout=(∃u∈users-usrs:chkdout(u,b))}
{bnd t: |usrs|}
do usrs≠{} → < S > od
{R: isout = (∃u∈users:chkdout(u,b))}
```

We now develop a statement that decreases the bound. We declare a local variable "usr" of type "user" and use the statement sequence "choose(usrs,usr); usrs:=usrs-usr" to remove an element from "usrs". Using item five of the proof rule we verify that this decreases "|usrs|", and it becomes the first part of the loop body.

Finally, we must develop a statement that restores the invariant after the previous command has invalidated it. There are two cases; therefore, the body contains an if statement. If "usr" has the book in question checked out ("chkdout(usr,b)"), then the result must be set to true; otherwise, nothing needs to be done.

Using item two of the proof rule we verify that the body of the loop maintains the invariant, and we have now produced the complete design shown in Figure 2. Since we ensured that all five items of the appropriate proof rule were satisfied as we constructed the loop, we have already proven it correct.

As we have just described it, the design process consists of a single level; the developer proceeds through a sequence of small, relatively independent steps to produce a final design. While we found this model adequate for performing walkthroughs, our

3

```
{Q: true}
var usrs : set(user) ;
var usr  : user        ;
isout,usrs:=false,users ;
{inv P:usrs⊆users ∧
      isout=(∃u∈users-usrs:chkdout(u,b))}
{bnd t: |usrs|}
do usrs≠{} →
        choose(usrs,usr); usrs:=usrs-usr;
        if  chkdout(usr,b) → isout:=true;
        ▯ ¬chkdout(usr,b) → skip;
        fi
od
{R: isout = (∃u∈users:chkdout(u,b))}
```

Figure 2. Completed *Isout* Design



Figure 3. Design Process as Simulated

attempts to automate the process based on it met with no success. Instead, our simulation program is based on a library of relatively large cliches representing solutions to common design problems.

## 3. Simulation Architecture

Figure 3 shows a pictorial representation of the design process implemented in our simulation program. It has two levels. At the lower level, the design derivation sub-process transforms formal specifications into verified designs using a library of cliches representing solutions to common programming problems. On the upper level, a cliche derivation sub-process uses strategies and proof rules to construct and verify cliches. These two sub-processes have significantly different complexities; cliche derivation is considerably more difficult than cliche application. Therefore, the portion of the process inside the dashed box is automated and the rest is performed by a human.

The input to the simulation is a pre- and post-condition specification for the unit to be constructed, as well as the library of pre-verified cliches. Each cliche has an applicability condition, as well as a rule for transforming specifications into more complete programs. The process simulation applies cliches until a complete design is produced or no cliches are applicable. The library of cliches is searched in a fixed order, with the simplest (least expensive to apply) cliches appearing first. Application of a cliche may generate sub-specifications for which a design must be created, and a simple backtracking scheme allows transformations to be undone if they do not lead to a complete
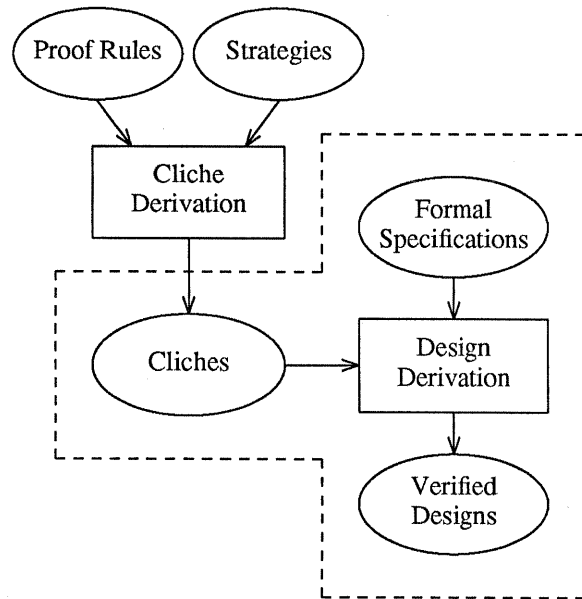
solution.

Since the correctness of a final design depends on the correctness of the cliches used in its derivation, each cliche must be proven to produce only designs that satisfy the corresponding specification. The advantage of our two level simulation architecture is that proofs are performed mostly at "compile" rather than "run" time. Cliche construction and verification is quite difficult, but is done only once for each cliche and performed by a human. On the other hand, cliche application is reasonably easy and is performed repeatedly by the machine.

The question of how a human performs the Gries/Dijkstra process is beyond the scope of this paper. It might be argued that even when someone performs a linear sequence of steps as described in section two, they are relying on a (possibly sub-conscious) cliche. On the other hand, some might say that the cliches are just a convenient way to store information that can be easily rederived when needed. As we understand it, the use of cliches is supported by work on the psychological aspects of programming [6,7,29,32].

4

We have constructed a simulation engine based on the architecture in Figure 3. The program was designed using guarded commands and its correctness rigorously verified [38]. It uses constructs that can be reasonably implemented in most programming languages. A prototype implementation has been written in Prolog that generates a complete design for several small examples including Kemmerer's Library Problem. The prototype follows the formal design very closely; in fact, the implementation can be generated from the design using methods similar to [35, 36]. The implementation is somewhat sketchy, especially the logic manipulation and theorem proving routines; however, it does demonstrate that the design is fundamentally correct.

For example, Figure 4 shows the code for the design process itself. The function "derive_design" takes a specification and if possible produces a complete design. In some cases it may not be able to produce a finished program, but it always preserves the "refines" relation. The body of "derive_design" uses the following sub-routines.

```
function optimize(d:dsgn)     : dsgn  ;
function derive_subs(d:dsgn): dsgn  ;
function complete(d:dsgn)     :boolean;
```

"Optimize" is a function that takes a design as input

---

```
function derive_design(s:spec):dsgn is

{Q: true}
var d : dsgn := s ; k : integer := 0 ;
{inv P: 0≤k≤|all_cliches| ∧ refines(s,d)}
{bnd t: |all_cliches|-k}
do k≠|all_cliches| ∧ ¬complete(d) →
    c,k := all_cliches[k],k+1 ;
    if c.pre(s) →
        d := optimize(
                derive_subs(c.apply(s)));
    ☐ ¬c.pre(s) → skip ;
    fi ;
od ;
derive_design := d ;
{R: refines(s,derive_design)}
end derive_design ;
```

Figure 4. Design Process Code

---

and returns a new one that has improved performance characteristics, while the function "derive_subs" takes a design and, if necessary, derives sub-designs to produce a complete program. Both of these functions preserve the "refines" relation. "Complete" is a function that returns true if all the unknowns in a design have been filled in; in other words, if the design is complete.

The body of "derive_design" consists of a single loop with an embedded conditional. Each iteration of the loop is concerned with a different cliche. If the current cliche is applicable to the specification, then it is applied and the result passed to "derive_subs" and then "optimize". If the cliche is not applicable then nothing is done. The loop terminates when a complete design is produced, or when all the cliches have been tried.

The architecture presented in this section is quite simple, but adequate for its purpose. The use of a library of cliches allows the design process to be separated into a difficult, possibly intuitive part performed by humans, and a simple, mechanistic part performed by a machine. The simplicity of the design derivation process implies that the power the overall process depends on the complexity of the cliches in its library. We have produced two distinct cliche libraries which differ significantly in their description of the design process.

## 4. Cliches - First Iteration

The number of cliches that can be used in the design process is literally infinite; however, a single example is sufficient to transmit the flavor of our first iteration. The real thrust of this effort was the construction of relatively powerful cliches that combine a number of steps from the walkthroughs into a single unit.

For example, Figure 5 shows a simplified representation of the "conditional_iteration_on_set" cliche. Application of this cliche can solve problems that require the use of a loop with an embedded conditional. In such cases, computation of the desired result involves processing each element of a set in turn. The post-condition for the cliche states that the result variable, "Var", is equal to the value of "Iop(Set,Cond)"; in other words, to the value of an iteration operator applied to a set with a certain condition.

The body of "conditional_iteration_on_set" declares two local variables. "Lset" is a set containing all the items still to be considered, while "Lvar" is the item currently being processed. "Lset" is initialized to "Set" and the result to the identity element, "Id". The loop interates over all the items in "Set". If the item in question satisfies "Cond" then "Var" is set to "Op(Var,Lvar)". The correct result has been calculated

```
cliche conditional_iteration_on_set is
    {Q}
    var Lset : set(Stype) ;
    var Lvar : Stype       ;
    Lset,Var := Set,Id      ;
    {inv P:Lset⊆Set ∧
          Var=Iop(Set-Lset,Cond) }
    {bnd t:|Lset|}
    do Lset≠{} →
        choose(Lset,Lvar) ;
        Lset:=Lset-Lvar    ;
        {Q1:Var=VAR}
        < S1 >( Var:inout Rtype ) ;
        {R1:¬Cond(Lvar) ∧ Var=VAR ∨
            Cond(Lvar) ∧
                Var=Op(VAR,Lvar) }
    od
    {R: Var = Iop(Set,Cond) } ;
if
    (Id,Op(Var,Lvar),Iop(Set,Cond))
            ∈ iop_table ;

end conditional_iteration_on_set ;
```

Figure 5. *Conditional_Iteration_on_Set* Cliche

when all the items have been considered.

The applicability test for this cliche is simplified by caching computationally expensive items. The cliche can be applied if the pre- and post-conditions unify and "iop_table" contains a tuple for the appropriate iteration operator. Each entry in the table contains an identity element as well as the operator used to update the result when an appropriate element is found.

A number of small cliches are also necessary to handle the details of design derivation. For example, the "simple_assignment" cliche generates (multiple) assignment statements. It states that a simultaneous assignment is correct if the pre-condition implies the post-condition with the left hand side of the assignment substituted for the right.

```
cliche simple_assignment is

    {Q} Var₁..Var_N := Soln₁..Soln_N {R}
if
    Q => R[[Var₁..Var_N / Soln₁..Soln_N]]

end simple_assignment ;
```

The "simple_if_then_else" cliche is somewhat more complex; it generates generates two branch if-then-else statements where one guard is the logical negation of the other.

```
cliche simple_if_then_else is
    {Q}
    if B1 → {Q ∧ B1} < S1 > {B1 ∧ E1}
    [] B2 → {Q ∧ B2} < S2 > {B2 ∧ E2}
    fi
    {R: B1 ∧ E1 ∨ B2 ∧ E2}
if
    is_negation(B1,B2) ;

end simple_if_then_else ;
```

All of the first iteration cliches are presented in more detail and verified in [38]. The main goal of this work was the construction of relatively powerful cliches, regardless of how they matched the process as described by the original authors [9, 10, 16, 17]. In line with our two level process model, cliche application was made as easy as possible, even at the cost of more effort in cliche construction. Although the framework we constructed is minimal, it does have enough power to duplicate some of the designs produced by a human.

### 4.1. Example Design Derivation

Let us reconsider the "isout" function presented in section two. We can see that the "conditional_iteration_on_set" cliche is applicable to the specification.

```
{Q: true}
< S >( isout:out boolean );
{R: isout = (∃u∈users:chkdout(u,b)) }

{Q}
< conditional_iteration_on_set >
{R: Var = Iop(Set,Cond) }
```

The specification and cliche unify with "Var" = "isout", "Iop" = "exists", "Set" = "users", and "Cond" = "chkdout(u,b)".

Figure 6 shows the result of applying the "conditional_iteration_on_set" cliche to the specification. The overall structure of the design is now evident. The loop iterates over all the users in the library. The variable "usrs" holds the users still to be considered, while "usr" holds the user currently being examined.

The loop body must still be completed before the design is finished. The "simple_if_then_else" cliche is applicable, and instantiation produces the following design for the loop body.

6

```
{Q: true}
var usrs : set(user) ;
var usr  : user        ;
isout,usrs:=false,users ;
{inv P:usrs⊆users ∧
      isout=(∃u∈users-usrs:chkdout(u,b))}
{bnd t: |usrs|}
do usrs≠{} →
    choose(usrs,usr); usrs:=usrs-usr ;
    {Q1: isout=ISOUT}
    < S1 >(isout:inout boolean);
    {R1:  chkdout(usr,b) ∧ isout=true ∨
         ¬chkdout(usr,b) ∧ isout=ISOUT}
od
{R: isout = (∃u∈users:chkdout(u,b))}
```

Figure 6. Instantiated Cliche

```
if  chkdout(usr,b) →
    {Q2:isout=ISOUT ∧ chkdout(usr,b)}
    < S2 >(isout:inout boolean);
    {R2:chkdout(usr,b) ∧ isout=true}
[] ¬chkdout(usr,b) →
    {Q3:isout=ISOUT ∧ ¬chkdout(usr,b)}
    < S3 >(isout:inout boolean);
    {R3:¬chkdout(usr,b) ∧ is_out=ISOUT}
fi
```

For each user, the loop body determines if they have the book in question checked out. If so, then "isout" is set to true, if not then nothing is done. The design of the loop body is completed by applying the "simple_assignment" cliche twice, thereby reproducing the hand derived design shown in Figure 2.

## 5. Cliches - Second Iteration

In the first iteration of our investigative process, we concentrated on constructing a system that could reproduce some of the designs created by a human; we did not devote a lot of effort to developing a process that closely matched the descriptions of the original authors [9,10,16,17], or the detailed activities of a human using pencil and paper. In the second iteration of our three part approach, we produced a new library of cliches that more closely reflect the method described by Gries in his book *The Science of Programming* [17].

The first iteration cliches combine a large number of steps from the walkthroughs into a single unit. In general, the second iteration cliches combine fewer steps than the first iteration cliches do; therefore, it may take more cliche applications to transform a specification into a completed design.

As an illustration, let us again consider the "isout" function. In this case, three second iteration cliches replace "conditional_iteration_on_set" from the first iteration. The "simple_replace_constant" cliche replaces a constant in a post-condition with a variable expression to create a loop invariant; the "decrease_then_restore" cliche decomposes a loop body into a statement which decreases the bound function and another that restores the invariant; and the "conditional_operation" cliche generates a loop body with an embedded conditional.

Figure 7 shows a simplified representation of the "simple_replace_constant" cliche. This cliche replaces a constant in a post-condition with a variable expression to create a loop invariant. Specifically, the cliche states that a loop with initialization "S0" and body "S1" is correct if: "C" is a constant in the post-condition, the tuple "(C,E,V,Rng)" is an element of the relation "mkvariable", and the loop invariant is equivalent to "Rng" conjoined with the result of substituting "E" for "C" in the post-condition. Here, "V" is a variable, "E" is an expression in "C" and "V", and "Rng" is a boolean expression restricting the range of "V".

```
cliche simple_replace_constant is
    {Q}
    var V ;
    < S0 > ;
    {inv P: Rng ∧ R[[ C / E ]]}
    {bnd t: |C|-|E|}
    do B: E≠C →
        {Q1: P ∧ B ∧ t=t1}
        < S1 > ;
        {R1: P ∧ t<t1}
    od
    {R}
if
    C ∈ constants(R) ∧
    (C,E,V,Rng) ∈ mkvariable ;

end simple_replace_constant ;
```

Figure 7. *Simple_Replace_Constant* Cliche

Figure 8 shows a simplified representation of the "decrease_then_restore" cliche. This cliche decomposes a loop body into a statement which decreases the bound function and another that restores the invariant. The applicability test for this cliche requires that the tuple "(T,V,E,S0,F,D)" is an element of the relation "decrease_bnd", and that "S1" does not modify any of the variables referenced in "T".

Here, "T" is the bound function for the loop; "V" is a variable that will be modified to decrease "T"; "S0" is a statement that modifies "V"; "E" is an expression reflecting the modifications to "V"; "F" is a formula stating additional facts concerning the modification; and "D" is a declaration of the iteration variable.

Figure 9 shows a simplified representation of the "conditional_operation" cliche. Application of this cliche generates a loop body with an embedded conditional. The post-condition for "conditional_operation" is the invariant for the enclosing loop. It states that "Lset" is a subset of "Set", and that "Var" is equal to the value of an iteration operator applied to the difference of "Set" and "Lset" with a certain condition.

The pre-condition for the cliche requires that the loop invariant hold with "LSET" substituted for "Lset", and that "Lset" be the result of removing "Lvar" from "LSET". Here, "Lset" is a set containing all the items still to be considered; "LSET" is the value of "Lset" from the previous iteration; and "Lvar" is the item currently being processed. "Conditional_operation"

```
cliche conditional_operation is

    {Q: LSET⊆Set ∧
        Var = Iop(Set-LSET,Cond) ∧
        Lset=LSET-Lvar ∧ Lvar∈LSET}
    {Q1: Var=VAR}
    < S1 >( Var:inout Rtype ) ;
    {R1: ( ¬Cond(Lvar) ∧ Var=VAR ∨
        Cond(Lvar) ∧ Var=Op(VAR,Lvar) )}
    {R: Lset⊆Set ∧
        Var = Iop(Set-Lset,Cond)}
if
    (Iop(Set,Cond),Op(Var,Lvar))
        ∈ iteration_ops ;

end conditional_operation ;
```

Figure 9. *Conditional_Operation* Cliche

can be applied to a specification if the pre- and post-conditions unify and the tuple "(Iop(Set,Cond), Op(Var,Lvar))" is an element of "iteration_ops".

The second iteration cliches are discussed in more detail and verified in [40]. They differ significantly from those developed in the first iteration of our investigative process. The second iteration cliches more closely match the process described in [17], and in general combine fewer of what were considered independent steps in our initial walkthroughs. Since they are in some sense "smaller", more applications may be required to produce a complete algorithm from a specification; however, they are still adequate to the task of reproducing the some of the designs created by a human.

### 5.1. Example Design Derivation

Let us again consider the "isout" function first presented in section two. The "simple_replace_constant" cliche is applicable to the initial specification. The constant "users" is selected from the post-condition; the "mkvariable" relation produces an expression "users-usrs" and range "usrs⊆users"; and substitution yields the following invariant.

```
cliche decrease_then_restore is

    {Q: P ∧ B ∧ T=t1}
    var D ;
    < S0 >
    {Q1: P[[ V / E ]] ∧ F}
    < S1 >
    {R1: P}
    {R: P ∧ T<t1}

if
    (T,V,E,S0,F,D) ∈ decrease_bnd ∧
    modify(S1) ∩ use(T) = ∅ ;

end decrease_then_restore ;
```

Figure 8. *Decrease_Then_Restore* Cliche

```
usrs⊆users ∧
isout = (∃u∈users-usrs:chkdout(u,b)))
```

Instantiation of the guard yields "users-usrs≠users", which simplifies to "usrs≠{}" under the assumption that "usrs" is a subset of "users". Similarly, the bound function instantiates to "|users|-|users-usrs|", which simplifies to "|usrs|". Therefore, application of "simple_replace_constant" produces the partial design shown in Figure 10.

Application of the "simple_assignment" cliche produces a loop initialization, and the "decrease_then_restore" cliche is applied to the specification of the loop body. The "decrease_bnd" relation produces a bound "|usrs|", variable "usrs", expression "USRS", statement "choose(usrs,usr); usrs:=usrs-usr", and formula "usrs=USRS-usr ∧ usr∈USRS". Therefore, instantiation of "decrease_then_restore" produces the following body for the loop.

```
choose(usrs,usr); usrs:=usrs-usr;
{Q1:USRS⊆users ∧
    isout=(∃u∈users-USRS:chkdout(u,b)) ∧
    usrs=USRS-usr ∧ usr∈USRS}
< S >(isout:inout boolean);
{R1:usrs⊆users ∧
    isout=(∃u∈users-usrs:chkdout(u,b))}
```

The "conditional_operation" cliche is applicable to the remaining specification, and instantiation

```
{Q: true}
var usrs : set(user) ;
< S0 >( isout:inout boolean ;
        usrs:inout set(user)) ;
{inv P:usrs⊆users ∧
    isout=(∃u∈users-usrs:chkdout(u,b))}
{bnd t: |usrs|}
do usrs≠{} →
    {Q1:usrs⊆users ∧
        isout=(∃u∈users-usrs:chkdout(u,b))
        ∧ usrs≠{} ∧ |usrs|=t1} ;
    < S1 >(isout:inout boolean);
    {R1:usrs⊆users ∧
        isout=(∃u∈users-usrs:chkdout(u,b))
        ∧ |usrs|<t1} ;
od
{R: isout = (∃u∈users:chkdout(u,b))}
```

Figure 10. Instantiated Cliche

produces the following.

```
choose(usrs,usr); usrs:=usrs-usr;
{Q1:isout=ISOUT}
< S >(isout:inout boolean);
{R1: chkdout(usr,b) ∧ isout=true ∨
    ¬chkdout(usr,b) ∧ isout=ISOUT}
```

The rest of the derivation proceeds exactly as described in section four; a single application of "simple_if_then_else" and two applications of "simple_assignment" reproduce the hand derived design shown in Figure 2.

## 6. Summary and Conclusions

We have been applying a three part, investigative approach to the formal design method described by Gries in his book *The Science of Programming* [17]. First, we walkthrough, in other words hand simulate, the process on a number of problems. This produces an increased understanding of the method as well as a suite of example designs. Second, we produce a program that simulates the design process discovered during the walkthroughs; ideally, it should be able to recreate the suite of designs previously produced. Third, we produce a process program that supports human application of the method.

We have currently completed two iterations of our three part approach on the Gries/Dijkstra process [37-41]; our understanding of the method has evolved significantly during the time we have spent studying it. After performing the initial walkthroughs, we viewed the design method as a linear sequence of small, relatively independent steps; however, our attempts to automate the process based on this model met with no success. Instead, both our simulation and process programs are based on a library of relatively large cliches representing solutions to common design problems.

Our simulation of the Gries/Dijkstra process has two levels. At the lower level, a design derivation sub-process transforms formal specifications into correct designs using a library of cliches. On the upper level, a cliche derivation sub-process uses strategies and proof rules to construct and verify cliches. The advantage of our two level architecture is that proofs are performed mostly at "compile" rather than "run" time. Cliche construction and verification is quite difficult, but is done only once for each cliche and performed by a human. On the other hand, cliche application is reasonably easy and is performed repeatedly by the machine.

We have constructed two distinct cliche libraries which differ significantly in their description of the design process. Our first iteration cliches combine a

number of what were considered distinct steps during the initial walkthroughs into a single unit. Our second iteration cliches combine several steps from the initial walkthroughs, but not nearly as many as the first iteration cliches do. They also better reflect the process described by Gries [17]. We believe our second iteration cliches are a significant step towards an accurate and detailed description of the method.

# 7. References

1. Barstow, D. R., "Domain-Specific Automatic Programming", *IEEE Trans. Software Eng. SE-11, 11* (Nov. 1985), 1321-1336.

2. Bernot, G. and M. C. Gaudel, "Software Testing Based on Formal Specifications: a Theory and a Tool", *Software Eng. J. 6*, 6 (Nov. 1991), 387.

3. Bjorner, D., "On The Use of Formal Methods in Software Development", *Proc. 9th Intl. Conf. Software Eng.*, 1987, 17-29.

4. Broy, M. and M. Wirsing, eds., *Methods of Programming: Selected Papers on the CIP-Project*, Springer-Verlag, New York, 1991.

5. Cunningham, H. C. and G. C. Roman, "A UNITY-Style Programming Logic for Shared Dataspace Programs", *IEEE Trans. Parallel Distributed Systems 1*, 3 (July 1990), 365.

6. Curtis, B., ed., *Tutorial: Human Factors in Software Development (Second Edition)*, IEEE Computer Society, Washington, D.C., 1985.

7. Davies, S. P., "The Role of Notation and Knowledge Representation in the Determination of Programming Strategy", *Cognitive Science 15* (1991), 547-572.

8. Delisle, N. and D. Garlan, "A Formal Specification of an Oscilloscope", *IEEE Software 7*, 5 (Sept. 1990), 29-36.

9. Dijkstra, E. W., "Guarded Commands, Nondeterminacy and Formal Derivation of Programs", *Comm. ACM 18*, 8 (Aug. 1975), 453-457.

10. Dijkstra, E. W., *A Discipline of Programming*, Prentice Hall, Englewood Cliffs, NJ, 1976.

11. Fagan, M. E., "Advances in Software Inspections", *IEEE Trans. Software Eng. SE-12*, 7 (July 1986), 744-751.

12. Fickas, S. and R. Helm, "Knowledge Representation and Reasoning in the Design of Composite Systems", *IEEE Trans. Software Eng. 18*, 6 (June 1992), 470-482.

13. Freedman, D. P. and G. M. Weinberg, *Handbook of Walkthroughs, Inspections and Technical Reviews*, Little, Brown and Company, Boston, 1982.

14. Futatsugi, K., J. Goguen, J. Meseguer and K. Okada, "Parameterized Programming in OBJ2", *Proc. 9th Intl. Conf. Software Eng.*, 1987, 51-60.

15. Goldberg, A. T., "Knowledge-Based Programming: A Survey of Program Design and Construction Techniques", *IEEE Trans. Software Eng. SE-12*, 7 (July 1986), 752-768.

16. Gries, D., "An Illustration of Current Ideas on the Derivation of Correctness Proofs and Correct Programs", *IEEE Trans. Software Eng. SE-2*, 4 (Dec. 1976), 238-244.

17. Gries, D., *The Science of Programming*, Springer-Verlag, New York, 1981.

18. *Proc. 7th Intl. Software Process Workshop*, IEEE Computer Society Press, Los Alamitos, CA, 1991.

19. Jones, C. B., *Systematic Software Development Using VDM*, Prentice-Hall, Englewood Cliffs, NJ, 1986.

20. Kant, E., "Understanding and Automating Algorithm Design", *IEEE Trans. Software Eng. SE-11, 11* (Nov. 1985), 1361-1374.

21. Linger, R. C., H. D. Mills and B. I. Witt, *Structured Programming: Theory and Practice*, Addison-Wesley, Reading, MA, 1979.

22. Loeckx, J. and K. Sieber, *The Foundations of Program Verification*, John Wiley & Sons, New York, 1984.

23. Lor, K. E. and D. M. Berry, "Automatic Synthesis of SARA Design Models From System Requirements", *IEEE Trans. Software Eng. 17*, 12 (Dec. 1991), 1229-1240.

24. Moriconi, M. S., "A Designer/Verifier's Assistant", *IEEE Trans. Software Eng. 5*, 4 (July 1979), 387-401.

25. Morzenti, A., D. Mandrioli and C. Ghezzi, "A Model Parametric Real-Time Logic", *ACM Trans. Programming Languages Systems 14*, 4 (Oct. 1992), 521.

26. Osterweil, L. J., "Software Processes Are Software Too", *Proc. 9th Intl. Conf. Software Eng.*, 1987, 2-13.

27. Potts, C. and G. Bruns, "Recording the Reasons for Design Decisions", *Proc. 10th Intl. Conf. Software Eng.*, April 1988, 418-427.

28. Rich, C. and R. C. Waters, eds., *Readings in Artificial Intelligence and Software Engineering*, Morgan Kaufman Publishers, Los Altos, CA, 1986.

29. Rist, R. S., "Knowledge Creation and Retrieval in Program Design: A Comparison of Novice and Experienced Student Programmers", *Human-Computer Interaction 6*, 1 (1991), 1-46.

30. Shaw, M., *Alphard: Form and Content*, Springer-Verlag, New York, 1981.

31. Smith, D. R., "KIDS: a Semiautomatic Program Development System", *IEEE Trans. Software Eng. 16*, 9 (Sept. 1990), 1024-1043.

32. Soloway, E. and K. Ehrlich, "Empirical Studies of Programming Knowledge", *IEEE Trans. Software Eng. SE-10*, 5 (1984), 595-609.

33. Sutton, S. M., D. Heimbigner and L. J. Osterweil, "Language Constructs for Managing Change in Process-Centered Environments", *Proc. 4th ACM SIGSOFT Symp. Software Development Environments*, Dec. 1990, 206-217.

34. Taylor, R. N., F. C. Belz, L. A. Clarke, L. J. Osterweil, R. W. Selby, J. C. Wileden, A. Wolf and M. Young, "Foundations for the Arcadia Environment Architecture", *Proc. Symp. Practical Software Development Environments*, 1988, 1-13.

35. Terwilliger, R. B. and R. H. Campbell, "An Early Report on ENCOMPASS", *Proc. 10th Intl. Conf. Software Eng.*, April 1988, 344-354.

36. Terwilliger, R. B. and R. H. Campbell, "PLEASE: Executable Specifications for Incremental Software Development", *J. Systems Software 10*, 2 (Sept. 1989), 97-112.

37. Terwilliger, R. B., "A Formal Specification and Verified Design for Kemmerer's Library Problem", Rprt. CU-CS-562-91, Dept. Comp. Sci., U. Colorado Boulder, Dec. 1991.

38. Terwilliger, R. B., "A Process Program for Gries/Dijkstra Design", Rprt. CU-CS-566-91, Dept. Comp. Sci., U. Colorado Boulder, Dec. 1991.

39. Terwilliger, R. B., "Towards Tools to Support the Gries/Dijkstra Design Process", Rprt. CU-CS-594-92, Dept. Comp. Sci., U. Colorado Boulder, May 1992.

40. Terwilliger, R. B., "A Second Simulation of the Gries/Dijkstra Design Process", Rprt. CU-CS-618-92, Dept. Comp. Sci., U. Colorado Boulder, Oct. 1992.

41. Terwilliger, R. B., "Simulating the Gries/Dijkstra Design Process", *Proc. 7th Knowledge-Based Software Engineering Conf.*, Sept. 1992, 144-153.

42. Weinberg, G. M. and D. P. Freedman, "Reviews, Walkthroughs, and Inspections", *IEEE Trans. Software Eng. SE-10, 1* (Jan. 1984), 68-72.

43. Wing, J. M., "A Study of 12 Specifications of the Library Problem", *IEEE Software 5*, 4 (July 1988), 66-76.

44. Woodcock, J. C. P., "Structuring Specifications in Z", *Software Eng. J. 4*, 1 (Jan 1989), 51.

45. Yourdon, E. N., *Structured Walkthroughs*, Yourdon Press, New York, 1989.