# The Design of Template Immune Networks: Path Immunity

Arthur M. Farley
Andrzej Proskurowski

CIS-TR-89-07
March 24, 1989

DEPARTMENT OF COMPUTER AND INFORMATION SCIENCE
UNIVERSITY OF OREGON

# The Design of Template Immune Networks: Path Immunity *

by

Arthur M. Farley
Andrzej Proskurowski

Department of Computer and Information Science
University of Oregon
Eugene, OR 97403

## ABSTRACT

A network is specified by a topology definition and a protocol definition. A network's topology, represented as a graph, defines its interconnection structure, while the protocol defines its operational behavior. A template is a connected graph. A topology G is immune to a set of templates T if G remains connected under removal of any imbedding of a single element of T. A network is template immune to a set of templates T if its topology is immune to T and its protocol guarantees that all operative sites can communicate in the presence of possible failures. A network is isolated template immune to a set of templates T if it is immune to multiple imbeddings of elements of T, where each imbedded template does not involve vertices that are neighbors of another imbedded template. We discuss networks that are isolated template immune to simple path templates of length k.

* Research supported in part by ONR contract #N00014-86.

1. Introduction

The robustness, or fault-tolerance, of a communication network is characterized by the set of communication tasks that can be completed in the presence of certain sets of failures in the network. In this report, we focus on site-to-site communication tasks in store-and-forward network architectures. Each site-to-site communication task is characterized by a *message*, an *originator*, being a site that creates the message, and a *recipient*, being the site to which the originator wishes to convey the message. In a store-and-forward network, messages are communicated by calls placed between adjacent sites. To be communicated between non-adjacent sites, a message must be forwarded by calls placed along a path connecting the two sites.

We model a communication network by a combination of topology and protocol specifications. A network's *topology* is represented by an undirected graph G =(V, E), whose vertices V correspond to sites of the network and whose edges E, represented as pairs of vertices (u, v), correspond to lines of the network. We refer to the graph representing a network's topology as the *network graph*. Two vertices are *neighbors* in a graph G if they are connected by an edge of G. A *path* (of length k) between two vertices $v_1$ and $v_k$ consists of a sequence of vertices $[v_1, ..., v_k]$, such that, for $0<i<k$, the pair $(v_i, v_{i+1})$ is an edge of the graph. We denote the general path on k vertices as $P_k$.

When evaluating the fault-tolerance of a network, we are concerned with connectedness properties of its network graph. Two vertices are *connected* in a graph G if there exists a path between them in G. A graph is *connected* if and only if every pair of vertices in G is connected. Removal of certain elements of a graph will disconnect the graph. A (vertex) *separator* of a graph G is a set of vertices, which, when removed from G, results in a subgraph of G that is not connected. A *minimal separator* S of G is a separator such that no proper subset of S is a separator of G.

A set of site failures in a network is modeled by removal of the corresponding vertices (and incident edges) from the network graph. In the case of a single site failure, the corresponding vertex and all edges incident to that vertex are removed from the graph. We refer to the graph remaining after removal of failed elements as the *operative graph*; the sites remaining are termed *operative sites*. Obviously, two operative sites can no longer successfully communicate if their associated vertices do not remain connected in a network's operative graph. However, two operative sites may not be able to communicate even when their corresponding vertices remain connected if the network's communication protocol is not sufficient to bypass extant failures.

A network's communication *protocol* is represented in terms of routing tables and calling procedures that are associated with each site of the network. In our networks of interest, communication is realized by a *call*, placed by one site, the *sender*, to another, neighboring site, the *receiver*. A *routing table* for a site indicates a sequence of one or more possible calls to its neighbors that the site may place for each other site as recipient. A *calling procedure* uses the routing table to determine which site(s), if any, to call with a message it has received or originated.

The site-to-site robustness, or fault-tolerance, of a communication network can be characterized by the set, or number, of pairs of operative sites that can still communicate in the presence of a given set of failures. We will say that a network is *immune* to a given set of failures if and only if all pairs of operative sites can still communicate in the remaining operative graph under the network's communication protocol.

## 2. Templated Failure Immunity

By *template*, we simply mean a connected graph. An imbedding of a template t in a graph G is a one-to-one association of the vertices of t with a subset of vertices in G, such that if there exists an edge between two vertices of t there is an edge between the associated vertices in G. There may be other edges in the subgraph of G induced by t; however, all edges of t must be in G. We say that the imbedding of a template in a graph G *covers* the associated elements of G. A *templated failure* is a failure of elements of a network that corresponds to an imbedding of a template in the network's graph. A templated failure is said to be based upon a set of templates T if the subgraph associated with the failure corresponds to one of the templates in T.

A graph G is *template immune* to a set T of templates if and only if it remains connected after introduction of any possible templated failure based upon T in G. The class of network graphs template immune to a set T of templates is referred to as *Immune(T)*. For example, if T consists solely of the graph U, where U denotes the trivial graph with one vertex and no edges, then Immune({U}) corresponds to the class of 2-connected graphs. Minimum examples of such graphs are cycle graphs, where, by minimum, we mean having the fewest edges for a given number of vertices. The *cycle graph* $C_n$ is the unique connected graph on n vertices such that every vertex has degree 2.

We can state the following lemma regarding graphs that are template immune to a set of templates T.

Lemma 2.1 A graph G is template immune to a set of templates T (i.e., G is in Immune(T)) if and only if no templated failure based upon T can cover a (minimal) separator of G.

A network N is template immune to a set T of templates if and only if its network graph is in Immune(T) and its communication protocol is such that every operative site can still communicate after introduction of a templated failure. Thus, that a network's graph is template immune is a necessary but not sufficient condition for the network to be template immune to a given set of templates.

An imbedding of elements from a set T of templates in a graph G is an *isolated imbedding* if it is an imbedding of templates (with possible duplication) from T such that no two of the imbedded templates have associated vertices that are neighbors in G. Around each failure there is a buffer of operative sites; as such, the failures are topologically isolated from each other. A graph G is *isolated template immune* to a set T of templates if and only if it remains connected after removal of any possible isolated imbedding of T in G. The class of graphs that are isolated template immune to a set of templates T will be denoted by *IsoImmune(T)*. We can state the following lemma regarding such classes of graphs.

Lemma 2.2 A graph G is isolated template immune to a set of templates T (i.e., G is in IsoImmune(T)) iff no isolated imbedding of templated failures based upon T in G can cover any (minimal) separator of G.

We have previously studied several classes of isolated template immune graphs [1,2]. IsoImmune({U}) is the class of graphs immune to removal of independent sets of vertices. In [1], we prove that 2-trees are in IsoImmune({U}). In [2], we broaden the characterization of such immune networks and define a class of networks in IsoImmune({$P_2$}).

The path graph $P_2$ can be generalized in several ways. One is to extend it in terms of the length of the path, looking at networks immune to $P_k$, for increasing k. Given the combinatorial explosion in number of connected graphs with increasing number of vertices, we will limit our attention to failure templates from the family of path graphs. Rather than consider immunity to single path failures, we will focus on the case of immunity to isolated instances of path template failures.

## 3. General Design Principles

If a graph is to be immune to isolated imbeddings of a particular set of templates, Lemma 2.2 gives a general, minimal condition that must be satisfied. Namely, no such imbedding can completely cover vertices of any separator in the graph. This condition is not particularly useful in designing the corresponding immune networks, as it simply restates the definition of immunity in terms of graph separators. What we want are several design principles upon which we can base our search for classes of graphs that are isolated immune to given sets of templates.

The notion of separator is obviously relevant, as removal of any separator disconnects the graph. If we could rephrase our condition in terms of sufficient conditions on the relationship between elements of T and separators of a graph, we would have principles to guide our search for feasible designs.

As noted above, when templated failures are isolated, there is a buffer zone of operative sites between any two failures. Therefore, as we imbed one failure, we guarantee that certain other vertices can not be included in other failures, when the failures are isolated. We will say that a given templated failure *immunizes* those vertices which must remain operative if failures are to be isolated in the network's graph. The immunization of nearby vertices suggests the following design principle P1, which is a sufficient condition for a network graph to have isolated immunity to a set of templates:

P1: Design a network so that for every separator, any templated failure covers only a proper subset of vertices of the separator and immunizes at least one other vertex of the separator.

We will refer to such separators as *self-immunizing separators*. If every separator of a graph G is self-immunizing for a given set of templates T, then G is in IsoImmune(T). One way to create self-immunizing separators is to consider network topologies in which every separator is (or, more generally, contains) a connected graph with more vertices than any template in T. Such graphs will be in IsoImmune(T).

Lemma 3.1 If every separator in a graph G contains a connected subgraph having more vertices than any element of a set of templates T, then G is in IsoImmune(T).

Proof: [By contradiction.] Assume there exists some isolated imbedding of elements of T that disconnects G. This implies that some separator S is completely covered by the imbedding. However, S contains a connected subgraph S' having more vertices than any template in T. Any templated failure that includes any part of S' cannot cover all of S' and must immunize at least one vertex of S'. In other words, every separator in G is self-immunizing. Thus, the imbedding of templates containing S can not be isolated, contradicting our assumption. []

If the graph U is in the set of templates T, then a templated failure only immunizes directly adjacent vertices. Any vertex not adjacent to a failure can fail by an imbedding of U. However, in those cases where the minimum diameter of the templates in T is greater than zero, non-adjacent vertices may be immunized by a given templated failure as well.

How can we construct (define) an infinite class of graphs that have a certain, desired structure in every separator? We need another design principle to guide our search. We employ a general technique, expressed as design principle P2, that is sufficient to generate such classes of graphs.

P2: Design a network by defining an immune base graph and a construction rule that adds a new immune subgraph by connections that create a self-immunizing separator.

The general technique we employ is one of iterative construction, growing desired graphs from a base graph by successively adding one or more vertices in a particular immune configuration, connecting them to a certain allowable subsets of vertices in the current graph. The base graph must be immune and each addition must maintain immunity by creating a self-immunizing separator between the two immune subgraphs. For example, a *k-tree* can be defined iteratively as either a k-complete graph (i.e., a graph on k vertices, such that every vertex is adjacent to every other vertex), or a k-tree K to which a new vertex v has been added by connecting v to a k-complete subgraph of K. The class of 2-trees are in IsoImmune({U}), by Lemma 3.1 and our design principles.

A related, recursive technique is to take two immune graphs and "glue" them together by identifying a certain subset of vertices from one graph with a subset in the other. The identified vertices become a separator that must satisfy conditions on immune structure. From this perspective, a 2-tree can be defined as being either a triangle or two 2-trees glued together by identifying a $P_2$ from each graph.

In our research, we are searching for *efficient* designs of networks immune to isolated imbeddings of certain sets of templates. As a measure of efficiency, we will use the number of edges (size) needed to realize immunity for a given number of vertices (order). There are a number of ways to evaluate the relative efficiency of a network design. One is to choose a reasonable, "benchmark" class (i.e., a relatively sparse, infinite family of immune graphs). We can evaluate the relative efficiency of an immune network design in terms of the improvement in size-to-order ratio between the new immune class and the benchmark graphs.

We will use the family of graphs known as k-trees as a benchmark for our designs. Every minimal separator in a k-tree is a k-complete graph. For a given set of templates T where the maximum number of vertices is k-1 in any template, a k-tree is immune to any isolated imbedding of elements of T, by Lemma 3.1. A k-tree has an size-to-order ratio of k. This provides us with a target ratio to improve upon in our search for efficient designs.

## 4. IsoImmune({U, $P_2$,... $P_k$}) Networks

In [1], we explored the definition of graphs immune to isolated instances of simple vertex or $P_2$ failures. Here, we consider networks where both of these failures may occur. Then we generalize to paths involving more vertices, defining topologies for networks immune to isolated occurences of these failures.

Let the*wheel graph* $W_n$ on n vertices be the graph having n-1 vertices connected in a cycle $C_{n-1}$, called the *rim*, and a single vertex, called the *hub*, connected to all vertices on the rim. By definition, $W_n$ has 2n-2 edges.

Theorem 4.1. The class of wheel graphs is exactly the class of minimum graphs in IsoImmune({U, $P_2$}).

Proof: It is easy to see that $W_n$ is in IsoImmuneI({U, $P_2$}). No single failure can disconnect the graph. If the hub is involved in a failure, only one (isolated) failure can occur. If the hub is not involved, the failure immunizes the hub; all operative vertices remain connected to each other indirectly through the hub.

To prove that $W_n$ is exactly the class of minimum size graphs in IsoImmune({U, $P_2$}), we first establish that the minimum size for a graph in this immune class is 2n-2, for graphs of order n. We see that $W_4$, the complete graph on 4 vertices, is the minimum size immune graph having order 4. Assume that G is a minimum order, immune graph (on n vertices) having fewer than 2n-2 edges. G must have a vertex u of degree 3, since G can have no vertex of degree 2, and, if all vertices were of degree at least 4, it would imply at least 2n edges. The neighborhood of u in G contains a path; without loss of generality, assume the path is <x,y,z>. We claim that the graph G' formed by removing vertex u from G and adding (if necessary) the edge <x, z> is also in IsoImmune({U, $P_2$}). This is so because any isolated imbedding of templates in G' that would disconnect G' would also disconnect G. Thus, G' is in IsoImmune({U, $P_2$}). G' has n-1 vertices and not more than 2n-2-2 = 2(n-1)-2 edges, as 3 edges were deleted from G and at most one added. Therefore, G was not of minimum order for an immune graph having fewer than 2n-2 edges. By contradiction, no such graph exists.

Now we show that wheels exactly constitute the class of minimum size graphs in IsoImmune({U, $P_2$}). Our argument is similar to that above for size. By inspection, $W_4$ (isomorphic to the complete graph on 4 vertices) is the only graph on four vertices that is in the class. Let k be the least k (where k > 4), such that $W_k$ is not the only graph having 2k-2 edges in IsoImmune({U, $P_2$}). Let G be one of these other graphs having k vertices. G must contain at least one vertex u of degree 3, as argued above. By Lemma 1, the neighborhood of u contains a $P_3$, <x, y, z>. Let us remove vertex u and add an edge between the two non-adjacent vertices x and z of the $P_3$. This reduces the number of vertices by 1 and number of edges by 2. The resultant graph G' is in IsoImmune({U, $P_2$}), as any disconnecting, isolated imbedding of v or $P_2$ in G' would also disconnect G. G' must be $W_{k-1}$, as k is the least k such that a minimum size member of IsoImmune({U, $P_2$}) is not be a wheel.

There are two cases to consider for the position of added edge <x,z>, since the triangle (x, y, z) resulting from the reduction always involves the hub (of $W_{k-1}$, since k > 4) and two adjacent vertices of degree 3. However, the added edge <x,z> could not be incident to the hub, as the original graph G would not be immune. Thus, the removed vertex u was on the rim of wheel graph $W_k$. Therefore, G was a wheel contradicting our definition of G. []

All that remains to complete the definition of networks immune to isolated instances of $P_2$ and single vertex failures is to define a communication protocol for sites of such a network. We call the following protocol $W_nP$.

If a site receives a call meant for it as recipient, the calling process is completed successfully. Otherwise, there are two classes of sites to consider, the hub site and the rim sites. When originating or forwarding a call, the hub site simply calls the ultimate recipient, as the hub is neighbor to all other sites. Regardless of whether the site is up or down, the calling process is complete and immune to failures.

If a rim site originates a message transfer, it first calls the hub. If the hub site is down, then the originator calls a neighbor on the rim, preferably the one closer to the ultimate recipient on the cycle of rim sites. If that neighbor is down and is not the ultimate recipient, the calling site calls its other neighbor on the rim, which must be operational. If forwarding a message, a rim site must have received the message from a neighbor on the rim (as the hub is down); it calls its other neighbor on the rim. If that site is down and it is not the ultimate recipient, the calling site calls the neighbor on the rim from which the message was received, thereby returning the message to traverse the cycle in the opposite direction.

Protocol $W_nP$ provides the immune behavior made possible by the wheel graph topology. Thus, we have the following theorem.

Theorem 4.2 Networks having wheel graphs as their topologies and using communication protocol $W_nP$ (as defined above) are immune to isolated $P_2$ and U failures.

One approach to the design of immune networks is to take graphs known to be immune to a special class and generalize those graphs to extend immunity to a larger classe of templated failures. We generalize our notion of wheel graph to define efficient networks immune to isolated imbedding of paths of length 1 (i.e., U) up to length k. Our benchmark class for such network graphs is the class of (k+1)-trees, with size to order ratio of k+1.

We define the class of multi-centered wheels, as follows. A *k-centered wheel* on n vertices $_kW_n$ (n> k+1) consists of a hub, being a set of k totally unconnected vertices, and a rim, being all other n-k vertices connected in a cycle, such that each vertex on the rim is also directly connected to every vertex of the hub. Figure 1 presents an illustration of a $_3W_{10}$ graph.

Theorem 4.3: The $_kW_n$ graph is in IsoImmune({U, $P_2$, ..., $P_{2k-2}$}), for n>3k-2 and k> 2.

Proof: We have k vertices in the hub and more than 2k-2 vertices on the rim. To disconnect a vertex on the rim, we must fail all hub vertices (plus at least two vertices on the rim) in one failure. To fail all vertices in the hub requires a path of length 2k-1. To disconnect a vertex in the hub requires that all vertices on the rim fail in one failure. To fail all vertices on the rim requires a path of length at least 2k-1. If a failure involves any vertex of the hub, all remaining vetices on the rim are immunized. If a failure involves a vertex on the rim, then all remaining vertices of the hub plus neighbors of the failure on the rim are immunized. As such, we see that a single failure from the set of available templates cannot disconnect the graph and immunizes sufficient other elements to guarantee immunity for the graph. []

The number of edges in a $_kW_n$ graph is n-k on the rim, plus (n-k)k connecting the rim to the hub, for a total of (k+1)(n-k). This indicates a size to order ratio of k+1 for networks in IsoImmune({U, $P_2$, ..., $P_{2k-2}$}), for k>2, which is significantly less than the size to order ratio of 2k-1 associated with our (2k-1)-tree, benchmark class.

Immune communication protocols for the $_kW_n$-based networks are straightforward generalizations of the $W_nP$ protocol defined previously. Each site on the rim now has a list of k hub sites to call (instead of only one). A site on the rim, when originating a message, calls the ultimate recipient first, if it is a neighbor of the recipient. Otherwise, it starts calling hub vertices until it succeeds (as it must by our proof of immunity above). A site on the hub forwards a message directly to its recipient. The only difficulty remaining for the protocol is a hub site originating a message for another hub site. The originator may have to call 2k-1 sites on the rim before finding one that is operative, which site can then call the recipient with the message. A call to an ultimate recipient, whether succeeding or failing, ends execution of the protocol. It is easy to see that this generalized protocol produces the immune communication behavior desired.

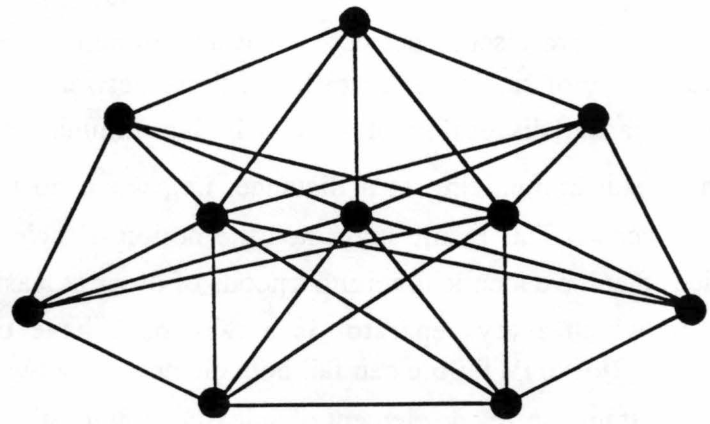

Figure 1. The ${}_3W_{10}$ graph.

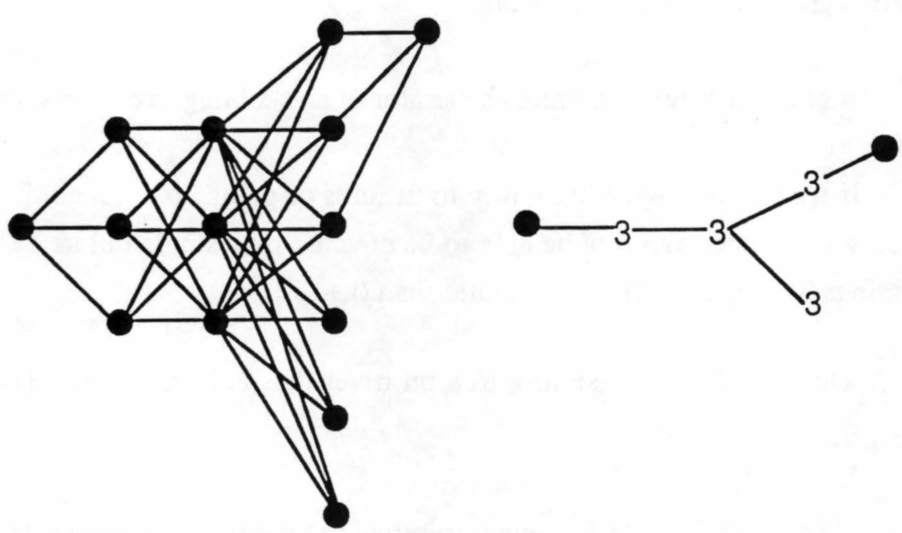


Figure 2. A 3-sibling tree.

## 5. IsoImmune($\{P_2, ..., P_k\}$) Networks

In this section we discuss networks that are immune to isolated path failures, without the possibility of single vertex failures, those networks in IsoImmune($\{P_2, ..., P_k\}$). From an earlier discussion of graphs in IsoImmune($\{P_2\}$), we found that a separator can be self-immunizing at a distance, i.e., when no pair of vertices of the separator are adjacent. This result suggested the notion of *twin pair,* being two non-adjacent vertices that have identical neighborhoods of order at least 2. Graphs known as *twin-trees*, in which every separator is a twin pair, have been shown to be in IsoImmune($\{P_2\}$) [2]. No $P_2$ failure can fail both elements of a twin pair and, in addition, each $P_2$ failure must immunize one element of at least one twin pair.

Our notion of twin pair can be generalized to that of an *s-sibling set* : a set of s independent (i.e., mutually non-adjacent) vertices having identical neighborhoods of order at least s. By an interative construction procedure, we can form tree-like graph structures, having s-sibling sets as separators.

An *s-sibling tree* on n vertices (for $n \geq 2s$) is either the complete bipartite graph on 2s vertices (for $n = 2s$) or is formed by connecting a new vertex v to an s-sibling-set of an s-sibling tree on n-1 vertices (for $n > 2s$). Each s-sibling tree on n vertices has s(n-s) edges. Figure 2 presents an example of a 3-sibling tree, including a representation highlighting the tree-like structure of such a graph. This recursive definition leads to an inductive proof of the following lemma.

Lemma 5.1 Every minimal separator of an s-sibling tree is an s-sibling set.

If we are to provide immunity to failures of paths up to length k, all vertices of a sibling set separator must not be able to be covered by a single failure. Thus, the number of siblings in a separator must be greater than (k+1)/2.

Theorem 5.2 An s-sibling tree on n vertices is in IsoImmune($\{P_2, ..., P_k\}$), for any s greater than (k+1)/2.

Proof: By Lemma 5.1, every minimal separator is an s-sibling set of order greater than (k+1)/2. Thus, no separator can be covered by a single $P_k$ failure. Each failure immunizes at least one vertex of any separator involved in the failure. □

The s-sibling trees are efficient designs in that they only require a size-to-order ratio of (k+1)/2, while our benchmark class of k+1-trees would require a ratio of k+1.

It remains to define immune communication protocols for s-sibling trees to complete our immune network design. In an s-sibling tree, either two sites are neighbors, and can call each other directly, or each vertex has a unique s-sibling set neighborhood between itself and any non-neighboring vertex. This is due to the tree-like structure of the sibling-set separators of the graphs. Therefore, the following, simply stated protocol provides the immune behavior desired: a site calls the ultimate recipient of a message if it is a neighbor of the recipient; otherwise, the site calls members of the unique, s-sibling set neighborhood lying between it and the ultimate recipient, until a call is successful.

## 6. Conclusion

In this paper we describe efficient designs for networks that are immune to isolated occurrences of path failures. Our designs include not only specifications of topology but also specifications of communication protocols that realize the immune performance. While efficient, most our designs are not minimum or not known to be minimum. Open questions include the determination of minimum designs for the cases considered here, other than the wheel graphs which were shown to be minimum elements of IsoImmune($\{U, P_2\}$).

## References

[1] Farley, A.M. "Networks immune to isolated failures," *Networks*, 11, 1981, 255-268.

[2] Farley, A.M. and Proskurowski, A. "Immunity to subgraph failures in communication networks", *Congressus Numerantium*, 62, 1988, p149-158.

[3] Klawe, M., Corneil, D., and Proskurowski, A. "Isomorphism testing in hookup classes", *SIAM J. on Algebraic and Discrete Methods*, 3(2), 1982, 260-274.