


Types for Lexically-Scoped Access Control

Tachio Terauchi

Alex Aiken

Jeffrey S. Foster



Report No. UCB/CSD-3-1282

October 2003

Computer Science Division (EECS)
University of California
Berkeley, California 94720

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE OCT 2003		2. REPORT TYPE		3. DATES COVERED 00-00-2003 to 00-00-2003	
4. TITLE AND SUBTITLE Types for Lexically-Scoped Access Control				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of California at Berkeley, Department of Electrical Engineering and Computer Sciences, Berkeley, CA, 94720				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT We develop a new system for defining and enforcing access control statically. In our system, key-pairs guard access to resources, and the association between key-pairs and resources can be changed at any program point (i.e., the binding is late). Our static system uses an ordering on lexically scoped abstract names to allow local access control policies to be enforced in other parts of a program. In particular this means that individual program components can locally refine access control policies and the policies will be respected by the entire program. The result is a system that can enforce, at compile time, a wide variety of useful, fine-grain access control patterns.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 34	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Types for Lexically-Scoped Access Control *

Tachio Terauchi

Jeffrey S. Foster

Alex Aiken

October 15, 2003

Abstract

We develop a new system for defining and enforcing access control statically. In our system, *key-pairs* guard access to *resources*, and the association between key-pairs and resources can be changed at any program point (i.e., the binding is late). Our static system uses an ordering on lexically scoped abstract names to allow local access control policies to be enforced in other parts of a program. In particular, this means that individual program components can locally refine access control policies and the policies will be respected by the entire program. The result is a system that can enforce, at compile time, a wide variety of useful, fine-grain access control patterns.

1 Introduction

In situations where a program P interacts with one or more untrusted program components U , a well-specified access control policy protects P 's resources from unwanted operations performed by U .

Systems such as Java and the Common Language Runtime (CLR) provide a mechanism for defining and enforcing fine-grain access control. In these systems, a programmer defines permissions denoting resources that pieces of code are allowed to access. For example, a Java class that conducts network transactions needs permission to access network sockets. A key part of *stack-based* access control is ensuring that if component C does not have permission to access resource v , then C cannot gain access to v by calling a function that has permission to access v . To enforce stack-based access control policies, the run-time system records the set of permissions S_1, S_2, \dots of each codes f_1, f_2, \dots on the execution stack. When a resource v is accessed, the run-time system permits the access only if $v \in \bigcap_i S_i$. That is, v can be used only if all functions on the stack have permission to access v . Stack-based access control systems usually also include a mechanism for gaining privileges, so that only a suffix of the stack is examined for access right.

In addition to a set of pre-defined permissions, Java and the CLR allow programmers to define new permissions. For example, in an airline ticketing program, a flight scheduling component may define new permissions used to protect internal resources such as flight time information from being modified by a ticketing agent component.

This paper presents a new stack-based access control system that is enforced statically, in contrast to the run-time enforcement in Java and the CLR. We begin with an example application of our system, slightly simplified for the purposes of introduction. Before giving the example, we give a very compressed overview of our approach.

In our system, permissions are expressed as *key-pairs* consisting of a *grant* key and a *limit* key. Each key-pair is associated with a set of resources, but the two keys are used for different operations on those resources. We have already mentioned that in stack-based access control each piece of code may limit access to a set of resources; this is the function of limit keys. The power to limit access is not sufficient, however. Something must also grant access to resources, at least initially, or no resources could ever be used; this is the function of grant keys. Splitting the management of access rights for a set of resources into grant and limit keys is fundamental to our design and, we believe, gives our system practical advantages (see Section 2).

*This research was supported in part by Subcontract no. PY-1099 to Stanford, Dept. of the Air Force prime contract no. F33615-00-C-1693.

A second novel feature of our system is *subkeys*. Subkeys are analogous to subclasses in an object-oriented language. If a limit (grant) key k_1 is a subkey of a limit (grant) key k_2 then k_1 guards access to a subset of the resources k_2 protects (see Section 2.5).

Finally, in our approach the programmer writes code to express access control policies using the following expressions for manipulating keys and resources:

- **newkey** generates a new pair of keys; **newkey** $< p$ generates new subkeys of the key-pair p . Initially, keys are not associated with any resources.
- **associate** e_1 with e_2 associates the key-pair e_1 with the resource e_2 .
- **limit** e_1 in e_2 limits e_2 's access rights to the resources associated with the limit key e_1 .
- **grant** e_1 in e_2 grants e_2 access rights to the resources associated with the grant key e_1 .
- The function **gKey**(p) (resp. **lKey**(p)) returns the grant (resp. limit) key of the key pair p .

Consider two components C and D that need to share access to a number of files. This is expressed in our system by creating a key-pair, associating it with the files, and then granting access to those files to both C and D:

```
shared_files = newkey;
...associate file_a with shared_files;
...associate file_b with shared_files;
[Component C]:
...
grant gKey(shared_files) in
  .../* accesses the shared files */
...
[Component D]:
...
grant gKey(shared_files) in
  .../* accesses the shared files */
```

Now assume that we wish to add to C a function f that is visible to other, less trusted, components, which may or may not have access to `shared_files`. We want only certain files in `shared_files` to be accessible through f . Our system is designed to support such refinements of access control policies with only local modifications to component C; in particular, neither component D nor any other component needs to be modified to implement these changes. We create a subkey-pair `some_files` of `shared_files` and associate it with a subset of files, say just `file_a` and `file_b`. Using `limit`, we ensure that callers of the function f can access at most the files associated with `some_files`; which files they can actually access depends on their access rights at the point where they call f .

```
[Component C]:
grant gKey(shared_files) in
...
some_files = newkey<shared_files;
...associate file_a with some_files;
...associate file_b with some_files;
...
f = λx.limit lKey(some_files) in
  .../* accesses the selected files */
```

Component D can call f because it has access to `shared_files`. But components without access to the selected files cannot.

In our view, our system provides two key features: late binding of keys and resources (i.e., a resource can be associated with a key after both the key and the resource exist) and the ability to create subkeys at any

point. As a result, these features allow late, local refinement of access control policies, as well as non-local enforcement of policies. In the example above, we have added a refinement of the original policy to function f in component C . This policy is local to C , meaning that the entire refinement is done in the component C , but it is enforced non-locally; i.e., a call to f in D respects the refined access control policy.

Our aim is to *statically* check such fine-grain access control policies. There are two benefits to static access control. One benefit is that it can find access violations early at compile time. The other benefit is efficiency; the run-time system need not to check for access violations.

Static checking requires our system to statically track three things: (a) the association between resources and keys; (b) the subkey hierarchy; and (c) which resources are used by expressions. For (a), we use abstract names, which are lexically scoped, to identify keys and to qualify resource types with their associated keys. It is important to note that only the static names of keys are lexically scoped; keys themselves are not lexically scoped and in fact can be passed as arguments, stored in data structures, etc. For (b), the static system tracks the *subnaming* hierarchy, which is the static analog of the subkey hierarchy. Finally, for (c) we develop a type and effect system that, for each expression e , assigns an effect set describing the set of resources accessed in the evaluation of e .

1.1 Contributions and Overview

This paper makes several contributions:

- Our access control system is static, i.e., the system checks for access violations at compile time. This is similar to some other recent research (e.g., [12]) but in stark contrast to common implementations of Java and CLR where violations are checked at run-time.
- Despite being static, the system is able to define and enforce many non-trivial access control patterns through its ability to reason about locally refined access control policies.
- The heart of our formal system is a novel, type-based must alias analysis, which may be of independent interest.
- We have a proof of soundness for the core subset of the system. Our system is sound in the presence of updatable references, higher-order functions, and concurrency.

Our system is type-based, relying only on standard techniques. One reason for choosing a type-based approach is compositional verification. Type systems explicitly express assumptions about the environment of each program fragment. Therefore program components may be checked separately under compatible environments, and then composed to form a well-typed program.

The rest of the paper proceeds as follows. Section 2 introduces our system informally. We introduce the important features of the system step-by-step, making improvements as we proceed. Materials up to and including Section 2.2 are sufficient to perform late, local refinement of access control policies. A small example is given in Section 2.4. To enforce local access policies non-locally, we introduce subkeys and subnaming in Section 2.5. Section 3 presents the static system formally and sketches a proof of soundness (the complete proof appears in Appendix A). Section 4 shows a few extensions to the system. Section 5 discusses our system's relation to must alias analysis. Section 6 discusses related work. Section 7 concludes.

2 Informal Presentation

This section informally develops the key ideas in our access control system. In the interests of clarity, we present some of the material introduced in Section 1 again, but in a more leisurely fashion and with more examples. The formal details of our system are presented in Section 3.

In a typical implementation of a dynamic stack-based access control system (such as Java's), permissions are objects containing strings, and the run-time system checks permissions by inspecting the strings before each access of associated resources. For example, permission to access flight information on flight 356 may be represented as a string "Flight Info : 356". Instead, we design access control into the static semantics of the language so that the type system can precisely track access control policies.

We introduce *key-pairs*, each consisting of a *limit key* and a *grant key*. Each key-pair is used to group resources. Limit keys are used to limit access to resources and grant keys are used to grant access to resources.

We consider every program value (e.g. a function, a pointer, etc) to be a resource possibly associated with some key-pair. Key-pairs are named in the static system using an infinite set `Names`. That is, each $\rho \in \text{Names}$ identifies a key-pair. While technically ρ is a static name for a run-time key-pair, for brevity we say “the key-pair ρ ” instead of “the key-pair identified by ρ ” when there can be no confusion.

Each type τ in our system consists of a name and a *raw type*. A resource of the raw type σ associated with the key-pair ρ is given the type $\rho \sigma$; we say that ρ *qualifies* σ in this type. For example, each pointer type is of the form $\rho \text{ref}(\tau)$.

The limit (resp. grant) key part of the key-pair ρ has the raw type *lkey* (ρ) (resp. *gkey* (ρ)). Keys are themselves program values. Therefore keys may also be under access control, and thus are associated with a key-pair. A limit (resp. grant) key ρ associated with the key-pair ρ_1 is given a type $\rho_1 \text{lkey}(\rho)$ (resp. $\rho_1 \text{gkey}(\rho)$).

Finally, we use $\perp \sigma$ to denote the type of a program value not associated with any key-pair. Such program values are not subject to access control, hence they may be accessed by any code.

2.1 Limit and Grant

To limit access rights, we introduce the syntax form

$$\text{limit } e_1, e_2, \dots, e_{n-1} \text{ in } e_n$$

where e_1, e_2, \dots, e_{n-1} are the limit keys and e_n is the code whose access rights are limited to the resources associated with the keys e_1, e_2, \dots, e_{n-1} .

Let `lkey` be the limit key of key-pair ρ . Assume that the function `h` and the pointer `t` are associated with `lkey`, and consider this definition:

$$\mathbf{g} = \lambda x. \text{limit } \text{lkey} \text{ in } x(\mathbf{h}, !\mathbf{t})$$

Since `g` dereferences `t`, `g` can be called only by code with access to `t`. For example, the following code should be rejected if the limit key `akey` is not associated with `t`. ($\lambda(x_1, \dots, x_n).e$ is a function of n arguments.)

$$\text{limit } \text{akey} \text{ in } \mathbf{g}(\lambda(x, y).y)$$

Moreover, because the access rights of the body of `g` are limited to resources associated with `lkey`, the function passed to `g` may only access resources associated with `lkey`. So if `lkey` is associated only with `h` and `t`, then `g` is forbidden from accessing any resource but `h` or `t`. For example, if `u` is a pointer not equal to `t` (and is not qualified with \perp)

$$\mathbf{g}(\lambda(x, y).!\mathbf{u})$$

is forbidden.

Now that we have seen how to limit access rights, we turn to granting access rights. We introduce the syntax form

$$\text{grant } e_1 \text{ in } e_2$$

where e_1 is a *grant key* and e_2 is the code granted access to the resources associated with key e_1 . Note that `grant` takes only a single expression as an argument, in contrast to `limit`, because multiple sequential `grants` can be used to grant a set of permissions, where multiple sequential `limits` will limit access to the intersection of access rights in the `limit` expressions.

Let `gkey` be the grant key part of the key-pair ρ . The code below grants itself access to the resources `h` and `t` in order to call `g`.

$$\text{grant } \text{gkey} \text{ in } \mathbf{g}(\lambda(x, y).x(y))$$

A **grant** is stronger than a **limit** because it enables unconditional access to the resources regardless of what access rights are available to the code’s context. In Java, the analog of **limit** is the permissions attached to traversed methods and the analog of **grant** is the initial set of permissions granted for an execution thread and additional accesses granted by the special form `doPrivileged`.

We use a type and effect system [7, 9] to statically check access control policies. For each expression, we assign an effect set conservatively approximating the expression’s access rights. Effect sets are finite sets of key-pair names. The type checking rules derive judgments of the form $\Gamma \vdash e : \tau; L$ (read “the expression e has the type τ in the type environment Γ and requires access rights L ”). For example, the type checking rule for pointer dereferences is:

$$\frac{\Gamma \vdash e : \rho \text{ ref}(\tau); L}{\Gamma \vdash ! e : \tau; L \cup \{\rho\}}$$

This rule says that the dereference of a pointer is well-typed if the access to the pointer is enabled.

The type checking rule for **limit** is more elaborate:

$$\frac{\text{For } 1 \leq i \leq n-1, \Gamma \vdash e_i : \rho'_i \text{ lkey}(\rho_i); L_i \quad \Gamma \vdash e_n : \tau_n; L_n \quad \Gamma \vdash L_n \subseteq \{\rho_1, \rho_2, \dots, \rho_{n-1}\}}{\Gamma \vdash \text{limit } e_1, e_2, \dots, e_{n-1} \text{ in } e_n : \tau_n; L_n \cup \bigcup_{1 \leq i \leq n-1} (L_i \cup \{\rho'_i\})}$$

The second line says that the body of the expression, e_n , only accesses the resources associated with the key-pairs $\rho_1, \rho_2, \dots, \rho_{n-1}$. In the conclusion, note the presence of L_n requires that the context have the right to access the key-pairs in L_n ; this prevents e_n from gaining more access rights than its context. On the first line, since each key e_i is itself a resource, it is associated with a key-pair ρ'_i (if it is not associated with any key-pair then $\rho'_i = \perp$). Moreover, evaluating each e_i may also involve resource accesses, approximated by L_i . Therefore, the context must have access rights to the resources associated with key-pairs identified by $\bigcup_{1 \leq i \leq n-1} (L_i \cup \{\rho'_i\})$, which is included in the overall effect.

The type checking rule for **grant** is simpler:

$$\frac{\Gamma \vdash e_1 : \rho_1 \text{ gkey}(\rho_2); L_1 \quad \Gamma \vdash e_2 : \tau; L_2 \cup \{\rho_2\}}{\Gamma \vdash \text{grant } e_1 \text{ in } e_2 : \tau; L_1 \cup L_2 \cup \{\rho_1\}}$$

The key e_1 is required to have a type of the form $\rho_1 \text{ gkey}(\rho_2)$, and e_2 is granted the access to the resources associated with the key-pair ρ_2 as indicated by the presence of $\{\rho_2\}$ in e_2 ’s effect. However, unlike with **limit**, the context need not have access to ρ_2 .

Grant keys are dangerous because they enable unconditional access rights regardless of the context. Hence programmers should prevent arbitrary code from using grant keys and gaining unwanted access rights. By splitting keys into key-pairs instead of using a single key for both limiting and granting, we have sought to make it easy to isolate and control the ability to grant access to resources. In particular, this design allows programmers to code in a style where each grant key is forgotten immediately after it is used to grant the access right to the code that generated the key.

Nevertheless, grant keys can be useful outside of this trivial style as long as they are used with discretion. As an example, consider system calls in the UNIX operating system. Application programs are usually prohibited from accessing system resources such as printers except through system calls. We can emulate the behavior of such a protocol using grant keys. Let all printers be associated with the grant key `k_printers`. Then the system function `print` can be written

```
λ(x,y).grant k_printers in {
  .../* code for outputting x to the printer y */ }
```

so that a program that does not have access to the printers may print by calling this function. ¹

¹The Java security model has a similar feature: `doPrivileged` enables a code to access more resources than are available to the code’s context.

2.2 Associating Resources with Key-Pairs

We next introduce the syntax forms **newkey** to generate new key-pairs and **associate** e_1 **with** e_2 to associate resources with key-pairs. The syntax form **newkey** $< e$, which generates new subkey-pairs, is introduced together with subnaming in Section 2.5.

We first describe **newkey**. We use existential types to guarantee that different key-pairs (particularly those generated by the same syntactic occurrence of **newkey**) are identified by different names.

$$\frac{}{\Gamma \vdash \mathbf{newkey} : \perp (\exists \rho. \perp \langle \perp \mathit{lkey}(\rho), \perp \mathit{gkey}(\rho) \rangle); \emptyset}$$

Raw tuple types are written $\langle \tau_1, \dots, \tau_n \rangle$. For extracting grant (resp. limit) keys from a pair, we define **lKey** (resp. **gKey**) as the projection function $\lambda x.(x.1)$ (resp. $\lambda x.(x.2)$).

Note that the type of a newly generated key-pair is qualified with \perp , i.e., by default the new key-pair is not associated with any key-pair and hence is available to any code. Other program values are also qualified with \perp when they are created, also making them available to any code. For example, the type checking rule for allocating a new store location (i.e. a new pointer) is

$$\frac{\Gamma \vdash e : \tau; L}{\Gamma \vdash \mathbf{ref} e : \perp \mathit{ref}(\tau); L}$$

In our system, each association between a key and a resource occurs after both the key and the resource are created, i.e., associations happen late. We use the syntax form

associate e_1 **with** e_2

to associate the resource e_1 with the key-pair with limit key e_2 .² The type checking rule for **associate** is shown below.

$$\frac{\Gamma \vdash e_1 : \rho_1 \sigma; L_1 \quad \Gamma \vdash e_2 : \rho_2 \mathit{lkey}(\rho_3); L_2}{\Gamma \vdash \mathbf{associate} e_1 \mathbf{with} e_2 : \rho_3 \sigma; L_1 \cup L_2 \cup \{\rho_1, \rho_2\}}$$

The **associate** construct is similar to *type casting*. That is, given a resource of the type $\rho_1 \sigma$, associating it with the key-pair ρ_3 has the effect of casting the type to $\rho_3 \sigma$.

In the rule for **associate**, because the limit key e_2 may itself be associated with some key-pair (identified by ρ_2), we assert that the context must hold the access right to ρ_2 . We assert that the context must also have access to the resource e_1 , and this is indicated by the inclusion of ρ_1 in the effect set of the conclusion. To see why this condition is necessary, consider the code below which might not be permitted to access the pointer **ptr** but nevertheless is able dereference it by generating a new key-pair, enabling, and associating **ptr** with it. (The form **open** $x = e$ unpacks the existential package e .)

```

open akeypair = newkey;
grant gKey(akeypair) in {
  ptr_casted = associate ptr with lKey(akeypair);
  ! ptr_casted; /* ptr accessed */
}

```

This code would type check if ρ_1 were missing from the effect set.

2.3 Preventing Resource Forging

Pointers, i.e. store locations, are naturally unforgeable in a strongly typed language that does not admit explicit pointer manipulations such as casting an integer to a memory address. In our system, keys are also unforgeable. But many other resources are forgeable. For example, any program can duplicate any pure function consisting only of variables and λ 's by simply having another instance of the function.

²We can also allow associating via grant keys and key-pairs themselves, as used in the example in Section 1. However, the ability to associate via limit keys is important, as it eliminates the need to expose grant keys (either alone or in key-pairs) for this purpose.

Our system does not automatically prevent such resource forging. This sounds unsafe because if a resource is forgeable then any code may re-create it and use the resource. Consider the file access control policy from Section 1. If file resource are strings containing the actual file name in the file system, then in any language with normal string operations it is trivial to create from scratch any file name with no access controls whatsoever.

Fortunately, such situations usually can be remedied easily. In the file access control example, we can define file resources to have an *abstract data type* ([10]) `File` defined in a program component `File Manager`. The representation of a `File` could be a string type, but this would not be visible outside of the `File Manager` component. In this design, `File Manager` would handle all file operations directly via system calls, and all other components would be required to go through `File Manager`. This can be made possible by using the access control mechanism to ensure that only `File Manager` has access to file-related system calls. When requesting some file for the first time, a component must ask `File Manager` for the appropriate `File` resource referring to the requested file. We also give `File Manager` a type signature such that, for example, calling the `File` writing function would have the effect of accessing the given `File` resource. Then `File Manager` component can ensure that `File` resources are unforgeable so as to allow access control on files by associating `File` resources with key-pairs.

In our system, we assume that a programmer takes necessary steps to make resources unforgeable if he wishes to have them under access control. For a naturally unforgeable resource like a pointer or a key, this can be as simple as associating the resource with some key-pair shortly after creating it and before making it available to other parts of the program. For other resources, use of abstract data types may be necessary, as seen in the `File` example.

2.4 Simple Example

The system described thus far is capable of expressing some fine-grain access control policies. Consider the following code which, for each node of the linked list `g`, applies the function `p` to the `item` field of the node.

```
t := g;
while (t ≠ nil) {
  p(t.item);
  t := (!t).next;
}
```

Let $\rho \neq \perp$. Let `g` have the record type `NodeType` defined below. (We take the liberty of using a C-like structure declaration for simplicity.)

```
NodeType =  $\rho$  {
   $\rho$  ItemType item;
   $\rho$  ref(NodeType) next;
};
```

(The first occurrence of ρ is not a binding name; it just qualifies the record type.)

We would like to use the system to limit `p` to accessing only the passed `t.item`. In particular, we want to prevent `p` from modifying the list structure of `g`, because the code is traversing it. We may do this by associating each `t.item` with a fresh key before calling `p`.

```
t := g;
while (t ≠ nil) {
  open akey = newkey;
  grant gKey(akey) in {
    theitem = associate t.item with lKey(akey);
    limit lKey(akey) in { p(theitem); };
    t := (!t).next;
  }
}
```

Note that the function `p` could be complicated, potentially calling other functions. Nevertheless, the type system is able to ensure that each function application only accesses the specified `t.item`, because the lexically scoped name (which appears when the new key-pair is opened) distinguishes each instance of `akey` and its late binding with the resource.

This example shows the system’s ability to declare late, local access control policies. That is, a programmer may store a collection of resources uniformly with the same type, and when there is a need for finer access control (e.g. at the level of individual resources in the collection) the key-associations are refined.

2.5 Subkeys and Subnaming

We now introduce subkeys, and their corresponding static representation using subnaming. These additions require a few changes to features explained earlier. To motivate the introduction of subkeys, we illustrate the problem that we would encounter if the system did not support them.

Consider the file access control example from Section 1. Recall that the key-pair `shared_files` is associated with the files shared between the program components `C` and `D`. Let ρ_{shared} identify `shared_files`. Below, the code for `C` is re-written; in particular, it uses `newkey` instead of `newkey<shared_files` to generate the key-pair `some_files`.

```
[Component C]:
open some_files = newkey;
...associate file_a with lKey(some_files);
...associate file_b with lKey(some_files);
...
f =  $\lambda x$ .limit lKey(some_files) in
    .../* accesses the selected files */
```

In the code above, `f` gets the type

$$\perp (\tau_1 \xrightarrow{\{\rho_{\text{some}}\}} \tau_2)$$

where ρ_{some} identifies the key-pair `some_files`. Here, the set above the arrow is the latent effect of the function, as usual in a type and effect system. Because ρ_{some} is lexically scoped in `C`’s context and therefore invisible in `D`’s context, the type system fails to type check `D`’s call to `f`.

The problem stems from the inability of the type system to understand local associations outside of the local context when the lexically scoped names are completely anonymous. This prevents non-local enforcement of locally refined policies.

Our system solves this problem by allowing the programmer to declare subset relations between resources via subkeys. The syntax form `newkey<e` generates a new immediate subkey-pair of an existing key-pair. As with `associate`, we let `e` be just the limit key part of the key-pair; again, this avoids the need to propagate the grant key just to associate resources with the key-pair. The type checking rule for `newkey<e` is

$$\frac{\Gamma \vdash e : \rho_1 \text{ lkey } (\rho_2); L}{\Gamma \vdash \text{newkey}<e : \perp (\exists \rho_3 < \rho_2. \perp (\perp \text{lkey } (\rho_3), \perp \text{gkey } (\rho_3))); L \cup \{\rho_1\}}$$

Here, in the bounded existential raw type $\exists \rho_3 < \rho_2. \dots$, the key-pair name ρ_3 is the bound name and ρ_2 is the upper-bound free in the raw type. The intuition behind this rule is that the new key-pair should be identified by some fresh name ρ_3 such that $\rho_3 < \rho_2$. The subnaming relation $<$ statically keeps track of the subkey hierarchy.

Recall that any resources associated with a key-pair are also associated with its superkey-pairs. Thus resources associated with a key-pair are a subset of its superkey-pair’s resources. Then, because the resources associated with the key-pair ρ_{some} are a subset of those associated with the key-pair ρ_{shared} , we can use `shared_files` as the superkey-pair when generating `some_files`. Here is the code from Section 1 again (using limit keys to generate and associate).

```

[Component C]:
  open some_files = newkey<lKey(shared_files);
  ...associate file_a with lKey(some_files);
  ...associate file_b with lKey(some_files);
  ...
  f =  $\lambda x$ .limit lKey(some_files) in
    .../* accesses the selected files */

```

Subnaming induces *subeffecting* and *subtyping*. In the example, the type system can reason that $\{\rho_{\text{some}}\} < \{\rho_{\text{shared}}\}$. We add the usual type checking rules so that a subeffect (resp. subtype) may be used where its supereffect (resp. supertype) is expected. Here we have

$$\perp (\tau_1 \xrightarrow{\{\rho_{\text{some}}\}} \tau_2) < \perp (\tau_1 \xrightarrow{\{\rho_{\text{shared}}\}} \tau_2)$$

and hence **f** can be called by a context having access to ρ_{shared} . Now D's call to **f** type checks.

Without subkeys and subnaming, the type system cannot track a key/resource association outside of the context of the code in which the association is made.³ The key observation that leads to subkeys and subnaming is that locally associated resources (i.e., a subkey) can be viewed as a local refinement of some larger collection of resources that is recognized (e.g., is associated with a superkey) outside the local context. In this way, subkeys and subnaming allow the programmer to define late, local associations that can be used non-locally.

To better understand the uses of subnaming, we present a larger example. Consider an airline ticketing program that contains two components: **FlightCtrl** having control of ticketing and seating information for all flights and **TicketAgt** representing a ticket agent. There is only one **FlightCtrl** but there could be more than instance of **TicketAgt**.

Suppose each component runs its own execution thread and that only read access to the seating and ticketing information is available to the **TicketAgt** thread, but **TicketAgt** may request **FlightCtrl** make modifications to the ticketing and seating arrangement. Let information for each flight be represented by the bounded existential record type

```

FlightInfo =  $\exists \rho_1 < \rho_{\text{all\_flights}}$  {
   $\perp$  lkey( $\rho_1$ ) key;
   $\rho_1$  TicketingInfo tickets;
   $\rho_1$  SeatingInfo seats;
   $\rho_2$  int flight_number;
   $\rho_2$  ref(FlightInfo) next;
};

```

The **FlightCtrl** thread is granted access to $\rho_{\text{all_flights}}$, so it may access the **FlightInfo** record of every flight.

Recognizing **TicketAgt**'s request, perhaps after some authentication, **FlightCtrl** responds by retrieving the flight information for the flight being requested:

```

[FlightCtrl Component]:
...
if (AuthTicketAgt(agent1, flight_num) {
  open finfo = getFlightInfo(flight_num);
...

```

The **TicketAgt** thread does not have write access to the **tickets** or **seatings** fields of **finfo**, so it cannot make the changes by itself. But it may create a function that does the job when called by **FlightCtrl**:

³To be precise, this statement is true up to absence of grant keys. In fact, without subnaming, the grant keys must be used to grant unconditional access to the non-local context when using associations non-locally. Therefore not having subnaming precludes many useful stack-based access control patterns.

```
[TicketAgt]:
...
f_agent = λx.
.../* Makes changes to
    finfo.tickets and finfo.seats */
```

FlightCtrl runs `f_agent` in its thread. But FlightCtrl, not fully trusting `f_agent`, uses the access control mechanism to check that `f_agent` is actually limited to modifying the seating and ticketing information of the specified flight:

```
[FlightCtrl Component]:
...
limit finfo.key in f_agent()
```

The type system guarantees that `f_agent` does not gain unwanted accesses (such as write access to flight information for other flights).

The subnaming relation $\rho_1 < \rho_{\text{all_flights}}$ implied in the type of `FlightInfo` is important. Existential quantification allows the program to operate on one `FlightInfo` at a time. Without subnaming, the bound name ρ_1 would become completely anonymous when the existential package is opened. Then, because this would mean that ρ_1 is a name unknown to `FlightCtrl`, `FlightCtrl` would be forbidden from accessing `finfo`, and hence would be unable to run `f_agent`. Subnaming allows anybody with access to $\rho_{\text{all_flights}}$ to perform access control at the level of individual `FlightInfo` instances.

We can take this example a step further. Suppose that `TicketAgt` also has its own set of resources associated with the key-pair ρ_{agt} . Assume that `f_agent` accesses some of these resources and that `FlightCtrl` has access rights to ρ_{agt} .

Let `ThirdPartyLib` be a third party library component that `TicketAgt` uses to, say, help itself calculate ticket costs and print out information on the terminal screen. To let `ThirdPartyLib` use some of its resources in a controlled manner, `TicketAgt` imposes access control on its own resources:

```
[TicketAgt]:
...
open lib_key = newkey;
...associate agt_resources.a with lKey(lib key);
...associate agt_resources.b with lKey(lib key);
```

Let ρ_{lib} identify the key-pairs `lib_key`. Now suppose that `f_agent` runs functions from `ThirdPartyLib`. The programmer wishes to limit the resources the functions can use, so the definition of `f_agent` would look like

```
[TicketAgt]:
...
f_agent = λx. {
...
  limit lKey(lib_key) in {
    .../* calls ThirdPartyLib's functions */ }
...
}
```

Coming back to `FlightCtrl`, it seems reasonable to assume that the `FlightCtrl` thread is still capable of running `f_agent` because it has accesses to `TicketAgt`'s resources as well as the flight information.

```
[FlightCtrl]:
...
/* agt_key is the key-pair identified by ρ_agt */
limit finfo.key, lKey(agt key) in f_agent()
```

$$\rho \in \text{Names} \cup \{\perp, \top\} \quad x \in \text{Vars} \quad m \in \text{Integers} \quad L \subseteq_{\text{finite}} \text{Names} \cup \{\perp, \top\}$$

raw types	$\sigma ::=$	$int \mid \langle \tau_1, \tau_2, \dots, \tau_n \rangle \mid ref(\tau) \mid \tau_1 \xrightarrow{L} \tau_2 \mid \forall \rho_1 < \rho_2. \tau \mid \exists \rho_1 < \rho_2. \tau \mid lkey(\rho) \mid gkey(\rho)$
types	$\tau ::=$	$\rho \sigma$
values	$v ::=$	$m \mid \text{Top} \mid x \mid \langle v_1, v_2, \dots, v_n \rangle \mid \lambda x : \tau. e \mid \Lambda \rho_1 < \rho_2. v \mid \text{pack } v \text{ as } \exists \rho_2 < \rho_3. \tau$
expressions	$e ::=$	$v \mid \langle e_1, e_2, \dots, e_n \rangle \mid e.i \mid e_1 e_2 \mid \text{ref } e \mid e_1 := e_2 \mid ! e \mid \text{spawn } e \mid e[\rho] \mid \text{pack } e \text{ as } \exists \rho_2 < \rho_3. \tau \mid \text{open } x = e_1 \text{ as } \tau \text{ in } e_2 \mid \text{limit } e_1, e_2, \dots, e_{n-1} \text{ in } e_n \mid \text{grant } e_1 \text{ in } e_2 \mid \text{newkey} < e \mid \text{associate } e_1 \text{ with } e_2$

Figure 1: Source language

But because we did not create `lib_key` as subkeys, the type system does not see the associations made in the `TicketAgt` component. Indeed, the type system must give `f_agent` the type $\perp (\tau_1 \xrightarrow{L} \tau_2)$ where $\rho_{\text{lib}} \in L$, but ρ_{lib} is not visible to `FlightCtrl`'s context.

Fortunately, subkeys provide an easy one-line fix.

```
[TicketAgt]:
...
open lib_key = newkey<lKey(agt_key); /* fixed */
...associate agt_resources.a with lKey(lib_key);
...associate agt_resources.b with lKey(lib_key);
```

Now the type system may give `f_agent` the type $\perp (\tau_1 \xrightarrow{L'} \tau_2)$ where L' contains ρ_{agt} instead of ρ_{lib} . Then the system is able to type check the call to `f_agent` in the `FlightCtrl` thread.

3 Formal System

We next formally present the core subset of our access control system and prove its soundness.

Figure 1 shows the source language, which consists of call-by-value first-class functions with primitives for imperative store operations and concurrency. The language is rather spare, missing a few bells and whistles seen in earlier sections for brevity (e.g., `while` and record types). These features can easily be added. We briefly explain the syntax.

As before, each type τ consists of a *name* and a *raw type*. The raw types are integers *int*, tuples $\langle \tau_1, \tau_2, \dots, \tau_n \rangle$, pointers *ref*(τ), functions $\tau_1 \xrightarrow{L} \tau_2$, bounded universal types $\forall \rho_1 < \rho_2. \tau$, bounded existential types $\exists \rho_1 < \rho_2. \tau$, limit keys *lkey*(ρ), and grant keys *gkey*(ρ). In bounded quantified types, the name on the left of $<$ is the bound name and the name on the right is the upper-bound free in the raw type.

A program value v is either an integer m , the constant limit key `Top`, a variable x , a tuple $\langle v_1, v_2, \dots, v_n \rangle$, a function $\lambda x : \tau. e$, a bounded polymorphic abstraction $\Lambda \rho_1 < \rho_2. v$, or a bounded existential package `pack` v `as` $\exists \rho_2 < \rho_3. \tau$.

A program expression e is either a value or one of the following syntax forms. The expression $\langle e_1, e_2, \dots, e_n \rangle$ creates a tuple. The expression $e.i$ projects the i th value of the tuple e . The expression $e_1 e_2$ applies the function e_1 to e_2 . The expression `ref` e allocates a new store location and initializes it to e . The expression $e_1 := e_2$ assigns e_2 to the store location e_1 . The expression `!` e dereferences the pointer (i.e. the store location) e . The expression `spawn` e spawns a new thread for the evaluation of e . The expression $e[\rho]$ instantiates the bounded universal abstraction e . The expression `pack` e `as` $\exists \rho_2 < \rho_3. \tau$ creates a bounded existential package of e . The expression `open` $x = e_1$ `as` τ `in` e_2 unpacks the bounded existential package e_1 to evaluate e_2 . Section 2 introduced the syntax forms `limit` e_1, e_2, \dots, e_{n-1} `in` e_n , `grant` e_1 `in` e_2 , `newkey` $< e$, and `associate` e_1 `with` e_2 .

Expressions and types are equivalent up to consistent renaming of bound names and variables.

$$\begin{array}{c}
\frac{}{\Gamma \vdash m : \perp \text{ int}; \emptyset} \text{ (Int)} \quad \frac{}{\Gamma \vdash \text{Top} : \perp \text{ lkey}(\top); \emptyset} \text{ (Top)} \quad \frac{x : \tau \in \Gamma}{\Gamma \vdash x : \tau; \emptyset} \text{ (Var)} \\
\\
\frac{\text{For } 1 \leq i \leq n, \Gamma \vdash e_i : \tau_i; L_i}{\Gamma \vdash \langle e_1, e_2, \dots, e_n \rangle : \perp \langle \tau_1, \tau_2, \dots, \tau_n \rangle; \bigcup_{1 \leq i \leq n} L_i} \text{ (Tuple)} \quad \frac{\Gamma \vdash e : \rho \langle \tau_1, \tau_2, \dots, \tau_n \rangle; L}{\Gamma \vdash e.i : \tau_i; L \cup \{\rho\}} \text{ (Proj)} \\
\\
\frac{\Gamma, x : \tau_1 \vdash e : \tau_2; L \quad \Gamma, x : \tau_1 \vdash \diamond}{\Gamma \vdash \lambda x : \tau_1. e : \perp (\tau_1 \xrightarrow{L} \tau_2); \emptyset} \text{ (Fun)} \quad \frac{\Gamma \vdash e_1 : \rho (\tau_1 \xrightarrow{L_1} \tau_2); L_2 \quad \Gamma \vdash e_2 : \tau_1; L_3}{\Gamma \vdash e_1 e_2 : \tau_2; L_1 \cup L_2 \cup L_3 \cup \{\rho\}} \text{ (App)} \\
\\
\frac{\Gamma \vdash e : \tau; L}{\Gamma \vdash \text{ref } e : \perp \text{ ref}(\tau); L} \text{ (Ref)} \quad \frac{\Gamma \vdash e_1 : \rho \text{ ref}(\tau); L_1 \quad \Gamma \vdash e_2 : \tau; L_2}{\Gamma \vdash e_1 := e_2 : \tau; L_1 \cup L_2 \cup \{\rho\}} \text{ (Assign)} \quad \frac{\Gamma \vdash e : \rho \text{ ref}(\tau); L}{\Gamma \vdash ! e : \tau; L \cup \{\rho\}} \text{ (Deref)} \\
\\
\frac{\Gamma \vdash e : \tau; \emptyset}{\Gamma \vdash \text{spawn } e : \perp \text{ int}; \emptyset} \text{ (Spawn)} \\
\\
\frac{\Gamma, \rho_1 < \rho_2 \vdash v : \tau; \emptyset \quad \Gamma, \rho_1 < \rho_2 \vdash \diamond}{\Gamma \vdash \Lambda \rho_1 < \rho_2. v : \perp (\forall \rho_1 < \rho_2. \tau); \emptyset} \text{ (Gen)} \quad \frac{\Gamma \vdash e : \rho_1 (\forall \rho_2 < \rho_3. \tau); L \quad \Gamma \vdash \rho_4 < \rho_3}{\Gamma \vdash e[\rho_4] : \tau[\rho_4/\rho_2]; L \cup \{\rho_1\}} \text{ (Inst)} \\
\\
\frac{\Gamma \vdash e : \tau[\rho_3/\rho_1]; L \quad \Gamma \vdash \rho_3 < \rho_2}{\Gamma \vdash \text{pack } e \text{ as } \exists \rho_1 < \rho_2. \tau : \perp (\exists \rho_1 < \rho_2. \tau); L} \text{ (Pack)} \quad \frac{\Gamma \vdash e_1 : \rho_1 (\exists \rho_2 < \rho_3. \tau_1); L_1 \quad \Gamma, \rho_2 < \rho_3, x : \tau_1 \vdash e_2 : \tau_2; L_2}{\Gamma \vdash \tau_2 \quad \Gamma \vdash L_2 \quad \Gamma, \rho_2 < \rho_3, x : \tau_1 \vdash \diamond} \text{ (Open)} \\
\\
\frac{\text{For } 1 \leq i \leq n-1, \Gamma \vdash e_i : \rho'_i \text{ lkey}(\rho_i); L_i \quad \Gamma \vdash e_n : \tau_n; L_n \quad \Gamma \vdash L_n < \{\rho_1, \rho_2, \dots, \rho_{n-1}\}}{\Gamma \vdash \text{limit } e_1, e_2, \dots, e_{n-1} \text{ in } e_n : \tau_n; L_n \cup \bigcup_{1 \leq i \leq n-1} (L_i \cup \{\rho'_i\})} \text{ (Limit)} \quad \frac{\Gamma \vdash e_1 : \rho_1 \text{ gkey}(\rho_2); L_1 \quad \Gamma \vdash e_2 : \tau; L_2 \cup \{\rho_2\}}{\Gamma \vdash \text{grant } e_1 \text{ in } e_2 : \tau; L_1 \cup L_2 \cup \{\rho_1\}} \text{ (Grant)} \\
\\
\frac{\Gamma \vdash e : \rho_1 \text{ lkey}(\rho_2); L}{\Gamma \vdash \text{newkey} < e : \perp (\exists \rho_3 < \rho_2. \perp \langle \perp \text{ lkey}(\rho_2), \perp \text{ gkey}(\rho_2) \rangle); L \cup \{\rho_1\}} \text{ (NewKey)} \\
\\
\frac{\Gamma \vdash e_1 : \rho_1 \sigma; L_1 \quad \Gamma \vdash e_2 : \rho_2 \text{ lkey}(\rho_3); L_2}{\Gamma \vdash \text{associate } e_1 \text{ with } e_2 : \rho_3 \sigma; L_1 \cup L_2 \cup \{\rho_1, \rho_2\}} \text{ (Associate)} \\
\\
\frac{\Gamma \vdash e : \tau_1; L_1 \quad \Gamma \vdash \tau_1 < \tau_2 \quad \Gamma \vdash L_1 < L_2}{\Gamma \vdash e : \tau_2; L_2} \text{ (Subsumption)}
\end{array}$$

Figure 2: Type checking rules: expressions

$$\begin{array}{c}
\frac{\Gamma \vdash \rho}{\Gamma \vdash \rho < \rho} \text{ (Sub Refl } \rho) \quad \frac{\Gamma \vdash \rho_1 < \rho_2 \quad \Gamma \vdash \rho_2 < \rho_3}{\Gamma \vdash \rho_1 < \rho_3} \text{ (Sub Trans } \rho) \quad \frac{\rho_1 < \rho_2 \in \Gamma}{\Gamma \vdash \rho_1 < \rho_2} \text{ (Sub } \rho) \quad \frac{\Gamma \vdash \rho}{\Gamma \vdash \perp < \rho} \text{ (Sub } \perp) \\
\frac{\Gamma \vdash \rho_1 < \rho_2 \quad \Gamma \vdash \sigma_1 < \sigma_2}{\Gamma \vdash \rho_1 \sigma_1 < \rho_2 \sigma_2} \text{ (Sub } \tau) \quad \frac{\Gamma \vdash \sigma}{\Gamma \vdash \sigma < \sigma} \text{ (Sub Refl } \sigma) \quad \frac{\Gamma \vdash \sigma_1 < \sigma_2 \quad \Gamma \vdash \sigma_2 < \sigma_3}{\Gamma \vdash \sigma_1 < \sigma_3} \text{ (Sub Trans } \sigma) \\
\frac{\text{For } 1 \leq i \leq n, \Gamma \vdash \tau_i < \tau_i'}{\Gamma \vdash \langle \tau_1, \tau_2, \dots, \tau_n \rangle < \langle \tau_1', \tau_2', \dots, \tau_n' \rangle} \text{ (Sub Tuple)} \quad \frac{\Gamma \vdash \tau_3 < \tau_1 \quad \Gamma \vdash \tau_2 < \tau_4 \quad \Gamma \vdash L_1 < L_2}{\Gamma \vdash \tau_1 \xrightarrow{L_1} \tau_2 < \tau_3 \xrightarrow{L_2} \tau_4} \text{ (Sub Fun)} \\
\frac{\Gamma \vdash \rho_3 < \rho_2 \quad \Gamma, \rho_1 < \rho_3 \vdash \tau_1 < \tau_2 \quad \Gamma, \rho_1 < \rho_3 \vdash \diamond}{\Gamma \vdash (\forall \rho_1 < \rho_2. \tau_1) < (\forall \rho_1 < \rho_3. \tau_2)} \text{ (Sub Univ)} \quad \frac{\Gamma \vdash \rho_2 < \rho_3 \quad \Gamma, \rho_1 < \rho_2 \vdash \tau_1 < \tau_2 \quad \Gamma, \rho_1 < \rho_2 \vdash \diamond}{\Gamma \vdash (\exists \rho_1 < \rho_2. \tau_1) < (\exists \rho_1 < \rho_3. \tau_2)} \\
\frac{\Gamma \vdash L}{\Gamma \vdash \emptyset < L} \text{ (Sub Effect } \emptyset) \quad \frac{\Gamma \vdash \rho_1 < \rho_2 \quad \Gamma \vdash L_1 < L_2}{\Gamma \vdash L_1 \cup \{\rho_1\} < L_2 \cup \{\rho_2\}} \text{ (Sub Effect } \{\rho\})
\end{array}$$

Figure 3: Type checking rules: subnaming, subtyping, subeffecting

$$\begin{array}{c}
\frac{}{\Gamma \vdash \perp} \text{ (Name } \perp) \quad \frac{}{\Gamma \vdash \top} \text{ (Name } \top) \quad \frac{\rho_1 < \rho_2 \in \Gamma}{\Gamma \vdash \rho_1} \text{ (Name } \rho) \\
\frac{\Gamma \vdash \rho \quad \Gamma \vdash \sigma}{\Gamma \vdash \rho \sigma} \text{ (Type)} \quad \frac{}{\Gamma \vdash \text{int}} \text{ (RType Int)} \quad \frac{\text{For } 1 \leq i \leq n, \Gamma \vdash \tau_i}{\Gamma \vdash \langle \tau_1, \tau_2, \dots, \tau_n \rangle} \text{ (RType Tuple)} \\
\frac{\Gamma \vdash \tau}{\Gamma \vdash \text{ref}(\tau)} \text{ (RType Ref)} \quad \frac{\Gamma \vdash \tau_1 \quad \Gamma \vdash \tau_2 \quad \Gamma \vdash L}{\Gamma \vdash \tau_1 \xrightarrow{L} \tau_2} \text{ (RType Fun)} \\
\frac{\Gamma, \rho_1 < \rho_2 \vdash \tau \quad \Gamma, \rho_1 < \rho_2 \vdash \diamond}{\Gamma \vdash \forall \rho_1 < \rho_2. \tau} \text{ (RType Univ)} \quad \frac{\Gamma, \rho_1 < \rho_2 \vdash \tau \quad \Gamma, \rho_1 < \rho_2 \vdash \diamond}{\Gamma \vdash \exists \rho_1 < \rho_2. \tau} \text{ (RType Exists)} \\
\frac{\Gamma \vdash \rho}{\Gamma \vdash \text{lkey}(\rho)} \text{ (RType LKey)} \quad \frac{\Gamma \vdash \rho}{\Gamma \vdash \text{gkey}(\rho)} \text{ (RType GKey)} \quad \frac{\text{For each } \rho \in L, \Gamma \vdash \rho}{\Gamma \vdash L} \text{ (Effect)} \\
\frac{}{\emptyset \vdash \diamond} \text{ (Env } \emptyset) \quad \frac{\Gamma \vdash \tau \quad \Gamma \vdash \diamond \quad x \notin \text{dom}(\Gamma)}{\Gamma, x : \tau \vdash \diamond} \text{ (Env } x) \quad \frac{\Gamma \vdash \rho_2 \quad \Gamma \vdash \diamond \quad \rho_1 \notin \text{dom}(\Gamma) \cup \{\perp, \top\}}{\Gamma, \rho_1 < \rho_2 \vdash \diamond} \text{ (Env } \rho)
\end{array}$$

Figure 4: Type checking rules: well-formed names, types, effects, environments

3.1 Static Semantics

Type environments are sequences of variable-to-type bindings and names annotated with a supertype.

$$\Gamma ::= \emptyset \mid \Gamma, x:\tau \mid \Gamma, \rho_1 < \rho_2$$

Figure 2 shows type checking rules for expressions. The meaning of $\Gamma \vdash e : \tau; L$ is that the effect set L conservatively approximates e 's access rights. A *program* is a closed expression. A program e is well-typed iff $\emptyset \vdash e : \tau; \emptyset$ for some τ , i.e., iff e can be type checked with no enabled access rights. Let us now discuss some of the important points in the type checking rules.

As explained in Section 2.5, (NewKey) assigns a bounded existential type to a newly created key-pair. The constant `Top` serves as the initial key. (Top) assigns `Top` the limit-key type $\perp \text{ lkey}(\top)$. The fact that `Top` is not a grant key is important, as otherwise any code would be able to access any resource by granting itself `Top`. The syntax form `newkey` is equivalent to `newkey < Top`.

Subnaming induces subtyping and subeffecting (see Figure 3), which are then used in two places: (Subsumption) and (Limit). (Subsumption) allows a subtype (resp. subeffect) to be used wherever its supertype (resp. supereffect) is expected. (Subsumption) formally captures the intention that whatever is associated with key-pairs are also associated with its super-key-pairs.

(Limit) replaces the rule introduced in Section 2 by using the subeffecting relation $<$ instead of the subset relation. This is needed to allow contexts with access rights limited to some key-pair to execute a code with access rights limited to its superkey-pair when the code accesses only what are allowed by the context. For example, consider the following function `g` which takes a resource and accesses it. We wish to limit `g`'s access rights to the key-pair `g_key` identified by ρ_g .

$$g = \Lambda \rho < \rho_g. \lambda x : \rho \sigma. \text{limit } \text{lKey}(g_key) \text{ in } \dots$$

The type system can give the `g` the raw type

$$\forall \rho < \rho_g. (\rho \sigma \xrightarrow{\{\rho\}} \tau)$$

for some σ and τ so that a caller whose access rights are locally refined with respect to ρ_g is able call `g` with its resource.

Figure 3 also shows that, as usual, subtyping is invariant under pointer types. Subtyping is also invariant for raw key types, i.e., $\text{lkey}(\rho_1) < \text{lkey}(\rho_2)$ then $\rho_1 = \rho_2$, and analogously for grant key types. It is easy to see that covariant keys would be unsound because limiting (resp. granting) ρ_1 should not somehow limit (resp. grant) ρ_2 when $\rho_1 < \rho_2$ and $\rho_1 \neq \rho_2$. To see that contravariance also fails, observe that associating ρ_1 should not somehow associate ρ_2 when $\rho_2 < \rho_1$ and $\rho_1 \neq \rho_2$.⁴

Section 2 discusses (Grant) and (Associate). Recall that values are created with no associated key-pair. Thus the rules for introductory forms (Int), (Tuple), (Fun), (Ref), (Gen), (Pack) and (NewKey) all qualify their values with \perp . The corresponding eliminatory forms (Proj), (App), (Assign), (Deref), (Inst), (Open), (Limit), (Grant), (NewKey) and (Associate) that use the values require an effect set consistent with the required access rights. Note that $\Gamma \vdash \perp < \rho$ for any $\Gamma \vdash \rho$.

As usual in a type and effect system, the latent effect of a function is recorded in its type at (Fun) and used at (App).

(Spawn) implies that a newly spawned thread starts without any access rights. We could alternatively have the spawned thread inherit the access rights of the spawner. In this case, we would add the rule

$$\frac{\Gamma \vdash e : \tau; L}{\Gamma \vdash \text{spawn } e : \perp \text{ int}; L} \text{ (Spawn alt)}$$

The only non-syntax directed rule is (Subsumption). Effect sets are finite and subtyping is structural, so type checking is decidable.

⁴We could relax invariance by assigning a pair of a covariant name and a contravariant name for each key type. Here we stick with invariance for simplicity.

3.2 Sketch of Type Soundness

We sketch a proof of soundness. A complete version of the proof appears in Appendix A.

The first step is to define a dynamic semantics rich enough to classify an access violation as a run-time error. This is only to prove soundness of the system; a real implementation should follow a straightforward dynamic semantics that ignores access control checks.

We introduce the set `KeyPairs` such that each element of `KeyPairs` denotes a key-pair. Each program value is now annotated with a key-pair, with the intended meaning that the value is associated with the key-pair. So a program value, say the function $\lambda x:\tau.e$, associated with the key-pair $k \in \text{KeyPairs}$ is

$$k \lambda x:\tau.e$$

We type check the annotated values by keeping track of names identifying $k \in \text{KeyPairs}$ in the type environment.

$$\Gamma ::= \dots \mid k:\rho$$

We use the special key-pairs `bot` and `top` to denote non-association and the `Top` key-pair, respectively.

The dynamic semantics needs to track enabled accesses for each context. Like in an eager stack-inspection semantics, we maintain exactly the enabled accesses in the syntax form `access A in e` where $A \subseteq_{\text{finite}} \text{KeyPairs} \cup \{\text{bot}, \text{top}\}$ denote the enabled accesses for the expression e . This approach makes the soundness proof easy. The following rule is used to type check `access A in e`.

$$\frac{\Gamma \vdash e : \tau; L \quad \Gamma \vdash L < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}}{\Gamma \vdash \text{access } A \text{ in } e : \tau; \emptyset}$$

`access A in e` is a run-time syntax form not available to the source program. Each `access A in e` appears as the result of the evaluation of `limit` or `grant`.

Evaluation is defined as a sequence of call-by-value small-step reductions of the form

$$(S, K, \langle e_1, e_2, \dots, e_n \rangle) \longrightarrow (S', K', \langle e'_1, e'_2, \dots, e'_{n'} \rangle)$$

where $e_1, e_2, \dots, e_n, e'_1, e'_2, \dots, e'_{n'}$ are the execution threads, S, S' are stores mapping store locations to program values, and K, K' are key-managers representing immediate sub-key-pair relation between key-pairs. So k_1 is a (immediate or non-immediate) sub-key-pair of k_2 iff $k_1 \mapsto k_2 \in K_{\perp}^*$ where K_{\perp}^* is the reflexive, transitive closure of K lifted with bottom. In this section, we restrict the reduction rules to the case for one execution thread to save space.

To make evaluation fail at access violations, the dynamic semantics checks the enabled accesses whenever a program value is used. For example, when calling a function associated with the key-pair k , the dynamic semantics checks whether k is enabled in the current set of access rights A .

$$\frac{k \in \text{Close}_K(A)}{(S, K, E[\text{access } A \text{ in } R[(k \lambda x:\tau.e) v]]) \longrightarrow (S, K, E[\text{access } A \text{ in } R[e[v/x]]])}$$

Close_K is a function such that $k \in \text{Close}_K(A)$ iff k is a sub-key-pair of some key-pair in A in the key-pair relation K . If the condition $k \in \text{Close}_K(A)$ is false, then this reduction cannot be taken, which implies that the evaluation for this thread gets stuck, indicating a run-time error.

The evaluation contexts E and R are used to find the redex and the set representing the enabled accesses for the redex. Intuitively, $E[\text{access } A \text{ in } R[e]]$ means that e is in an evaluation context and A is the set of its enabled accesses.

We use the special constant value `err` to signal a run-time error. Whenever an execution thread gets stuck, i.e., when it reduces to a irreducible non-value, it immediately reduces to `err`. Note that `err` has no type.

As mentioned before, `limit` and `grant` introduce occurrences of `access A in e`. The new access rights B after `limit` is reduced is the intersection of the previous access rights A with those specified in `limit`.

(Here $\mathbf{lkey}(k)$ denotes the limit key value (unannotated) of the key-pair k .)

$$\frac{\begin{array}{c} \{k'_1, k'_2, \dots, k'_n\} \subseteq \text{Close}_K(A) \\ B = \text{Close}_K(A) \cap \text{Close}_K(\{k_1, \dots, k_n\}) \end{array}}{(S, K, E[\text{access } A \text{ in } R[\text{limit } k'_1 \mathbf{lkey}(k_1), \dots, k'_n \mathbf{lkey}(k_n) \text{ in } e]]) \longrightarrow (S, K, E[\text{access } A \text{ in } R[\text{access } B \text{ in } e]])}$$

In contrast, the new access rights B after **grant** is reduced is the union of the previous access rights A with those specified in **grant**. (Here $\mathbf{gkey}(k)$ denotes the grant key value (unannotated) of the key-pair k .)

$$\frac{\begin{array}{c} k_1 \in \text{Close}_K(A) \\ B = A \cup \{k_2\} \end{array}}{(S, K, E[\text{access } A \text{ in } R[\text{grant } k_1 \mathbf{gkey}(k_2) \text{ in } e]]) \longrightarrow (S, K, E[\text{access } A \text{ in } R[\text{access } B \text{ in } e]])}$$

The reduction for **associate** is relatively simple. Here r is an unannotated program value. The reduction changes r 's annotation from k_1 to k_3 .

$$\frac{\{k_1, k_2\} \subseteq \text{Close}_K(A)}{(S, K, E[\text{access } A \text{ in } R[\text{associate } k_1 \text{ } r \text{ with } k_2 \mathbf{lkey}(k_3)])]) \longrightarrow (S, K, E[\text{access } A \text{ in } R[k_3 \text{ } r]])}$$

The reduction for **newkey** creates a new key-pair k_3 as the immediate sub-key-pair of k_2 .

$$\frac{\begin{array}{c} k_1 \in \text{Close}_K(A) \quad k_3 \notin \text{dom}(K) \cup \{\mathbf{bot}, \mathbf{top}\} \end{array}}{(S, K, E[\text{access } A \text{ in } R[\text{newkey} < k_1 \mathbf{lkey}(k_2)])]) \longrightarrow (S, K \cup \{k_3 \mapsto k_2\}, E[\text{access } A \text{ in } R[v]])}$$

where v is the existential package containing the new key-pair value (not shown).

Appendix A.1 shows the complete dynamic semantics.

Recall that we have intentionally enriched the dynamic semantics with explicit key-pairs just so that we could identify access violations. A real implementation of the run-time system does not need key-pairs because our type system checks for access violations at compile time.

The next step is to define a well-typed program state.

Definition 1 $\Gamma \vdash (S, K, \langle e_1, e_2, \dots, e_n \rangle)$ (read “the program state $(S, K, \langle e_1, e_2, \dots, e_n \rangle)$ is well-typed under Γ ”) iff

- (1) $\Gamma \vdash \diamond$.
- (2) For each $1 \leq i \leq n$, $\Gamma \vdash e_i : \tau_i; \emptyset$ for some τ_i .
- (3) $\ell \in \text{dom}(S)$ iff $\ell \in \text{dom}(\Gamma)$.
- (4) $k \in \text{dom}(K)$ iff $k \in \text{dom}(\Gamma)$.
- (5) $x \notin \text{dom}(\Gamma)$.
- (6) If $\Gamma \vdash \ell : \tau$ and $S(\ell) = v$ then $\Gamma \vdash v : \tau; \emptyset$.
- (7) Suppose $\Gamma \vdash k_1 : \rho_1$ and $\Gamma \vdash k_2 : \rho_2$. Then $\Gamma \vdash \rho_1 < \rho_2$ iff $k_1 \mapsto k_2 \in K_{\perp}^*$.

The first condition says that Γ must be well-formed. The second condition says that each execution thread is well-typed under Γ with effect set \emptyset . The third and the fourth conditions say that Γ and S, K must have matching store locations and key-pairs. The fifth condition says that Γ must not contain variables. The sixth condition expresses the usual well-typed store condition. The last condition says that the key-manager K contains exactly the sub-key-pair relationship implied by the type environment Γ .

Recall that **err** is non-tyable. The soundness proof then reduces to showing the following theorem.

Theorem 1 (Subject Reduction) *If $\Gamma \vdash (S, K, \vec{e})$ and $(S, K, \vec{e}) \longrightarrow (S', K', \vec{e}')$, then there is Γ' such that $\Gamma' \vdash (S', K', \vec{e}')$.*

The theorem is proved by induction on the type derivation and case analysis on reduction kinds. The immediate corollary is that a well-typed program does not cause access violations.

4 Extensions

We next discuss a few possible extensions to the system.

4.1 Run-Time Checks

Our system is powerful enough to statically enforce many fine-grain access control policies. But the programmer may still encounter situations where the static checking feels overly restrictive. We may extend our system with run-time checks.

$$\frac{\Gamma \vdash e_1 : \rho_1 \text{ lkey}(\rho_2); L_1 \quad \Gamma \vdash e_2 : \tau; L_2 \cup \{\rho_2\} \quad \Gamma \vdash e_3 : \tau; L_2}{\Gamma \vdash \text{have-access } e_1 ? e_2 e_3 : \tau; L_1 \cup L_2}$$

The syntax form `have-access $e_1 ? e_2 e_3$` checks whether the current context has e_1 enabled, and if so evaluates e_2 or else evaluates e_3 .

The disadvantage of this extension is that the run-time system now must track key values and enabled access rights. Without this extension, our system is completely static, and therefore a program with access control has no run-time overhead over an equivalent program without access control. Nevertheless, even with this extension, the run-time system only needs to check enabled access rights at `have-key`'s. Hence the overhead should be minimal.

4.2 Effect Kinds

In our airline ticketing example, we assumed that the `TicketAgt` thread has read access but not write access to the flight information. Strictly speaking, the system we described cannot distinguish different kinds of effects on the same resource.

It is easy to extend the system with effect kind constants for the primitives in the language (e.g., distinguishing the effect of dereferencing from the effect of assigning or associating) and programmer-defined effect kind variables (e.g., the effect kinds `readSeating`, `editSeating` for `SeatingInfo`). Effect kinds qualify effects, i.e., each effect is now of the form $\eta \rho$ where η is an effect kind. Then the programmer may, for example, limit e 's access to only reading a selected `SeatingInfo` by

`limit readSeating(theseating_key) in e`

and the type system checks that e 's effect set is a subeffect of $\{\text{readSeating } \rho\}$, where ρ identifies `theseating_key`.

4.3 Type Variables, Effect Variables, and More

We restricted the formal system to just bounded quantified types over names to make the presentation concise. But because the system relies only on standard type-system techniques, it is straightforward to extend the system with type variables and effect variables so as to admit bounded quantification over them.

Existentially quantified type variables naturally encode abstract data types [10]. We may also want a native support for parameterized recursive data types. A practical implementation may be to add recursively-defined named data types seen in languages like ML.

5 Relation to Must Alias Analysis

We briefly discuss the relation between our access control system and must alias analysis. In order to express locally refined access control policies, our access control checks constraints of the form “resources used by the expression e must-alias one of the resources in the set A ,” where A may be defined locally. The first step toward this goal is generation of lexically scoped names and late binding of names with program values. Lexically scoped names allow the creation of new names distinguishable from other names, and late binding allows existing program values to be must-aliased with these names. Together, they support creation of fine must-aliasing relations. Note that relations are not lexically scoped; they may be transported outside of a lexical scope by existential or universal name quantification.

However, there is a problem when these names are left disjoint. The problem occurs when the names collide (with other names or with their lexical boundaries), for example, when the type system needs to equate two types differing in names. The problem becomes more severe in a setting like stack-based access control because names propagate and collide not only through types of expressions but also through their effects. This is the problem subnaming solves. With subnaming, names are no longer disjoint. So the type system is able to make compromises by using supernames when possible. This has the effect of “downgrading” a fine must-aliasing relation to a coarser one.

6 Related Work

Stack-based access control is used in Java [8] and the CLR [4]. Pottier et al [12] present a type system that can statically check a subset of Java’s access control policies. Their system is faithful to the Java security model and hence conservatively approximates Java’s dynamic access control. It is not our goal to match Java’s access control system, but we believe it is possible to encode Java’s mechanism in a way so that access policies enforceable by Pottier et al’s system can also be enforced by ours. We also believe that our system is more expressive, because of its ability to locally refine access control policies.

The implementation of run-time systems for stack-based dynamic access control is a research subject in itself. In addition to direct inspection of the execution stack at run-time, which is done by Java and the CLR, several other techniques have been proposed. One approach is security-passing style where enabled access rights are explicitly passed as function arguments [15]. Also, an implementation technique based on code instrumentation has been proposed [13]. We believe that, should the need for run-time checks arise in our system (e.g., as discussed in Section 4.1), we can add such checks using any of these techniques. Besides efficient implementation, studying dynamic semantics has led to a deeper understanding of the stack-based access control mechanism [15, 6]. It is interesting to note that Pottier et al [12] first develop their static system on the security-passing-style converted language, then translate back to obtain a corresponding system for the source language.

There are a variety of ways to perform access control besides the stack-based approach. Abadi et al [1] present a system based on examining execution history. Their system is a generalization of the stack-based mechanism where the access rights at a point in time are determined not only by the access rights encoded in the stack but also by those of functions that have already returned. Using our syntax, their system roughly can be explained as splitting `limit` into `start_limit` and `end_limit` and analogously for `grant`. Statically checking history-based policies in a convincing manner would require reasoning about the sequential behavior of programs, and it is unclear whether we would be able to adapt our system to statically checking history-based access control policies.

Researchers have also proposed systems for statically enforcing generic security policies. These systems include Hoare style proof-systems [11, 3] and a higher-order type system [14]. One of their goals is a portable framework that is less tied to the source language. It is difficult to compare our system with these as the problem domains are so different.

Our system is an application of type and effect systems [7, 9]. Type and effect systems are widely used for checking non-standard program properties; the work of Pottier et al also uses a type and effect system [12].

Finally, there is a connection to work on aliasing control, particularly previous work on the programming construct `restrict`, originally studied in [5] and extended in [2]. Briefly, `restrict` is a mechanism for

specifying and checking local non-aliasing of store locations. To this end, abstract location names are used to approximate store locations so that location names distinguish local aliases of store locations from other aliases. Beyond what has been previously published on this subject, we noticed that, instead of using incomparable location names, we could use subnaming so that each local alias of a set of store locations S is a subtype of the type of all aliases of S . This captures the notion that each local alias is a refinement of the set of all aliases. Our use of subnaming for refinement of access control policies exploits the same idea.

7 Conclusions

We have presented a new system for defining and enforcing fine-grain access control. The system is simple, relying only on standard techniques from the types literature, yet powerful enough to express many non-trivial access control policies statically. The system draws its power from two key features: late and locally defined access control policies and non-local enforcement of the policies. We make the former possible through late binding between resources and key-pairs identified by lexically scoped names. To make the latter possible, we use subnaming.

At the core, our use of subnaming is a kind of local must alias analysis. This technique itself appears to be applicable to a variety of problems in addition to access control.

References

- [1] M. Abadi and C. Fournet. Access Control based on Execution History. In *Proceedings of the tenth Annual Network and Distributed System Security Symposium*, pages 107–121, Feb. 2003.
- [2] A. Aiken, J. S. Foster, J. Kodumal, and T. Terauchi. Checking and Inferring Local Non-Aliasing. In *Proceedings of the 2003 ACM SIGPLAN Conference on Programming Language Design and Implementation*, San Diego, California, June 2003.
- [3] C. Colby, K. Crary, R. Harper, P. Lee, and F. Pfenning. Automated techniques for provably safe mobile code. *Theoretical Computer Science*, 290:1175–1199, 2003. Special issue on *Dependable Computing*.
- [4] ECMA. *Standard ECMA-335: Common Language Infrastructure*, 2002. <http://www.ecma-international.org/publications/standards/ecma-335.htm>.
- [5] J. S. Foster and A. Aiken. Checking Programmer-Specified Non-Aliasing. Technical Report UCB//CSD-01-1160, University of California, Berkeley, Oct. 2001.
- [6] C. Fournet and A. G. Gordon. Stack Inspection: Theory and Variants. In *Proceedings of the 29th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 307–318, Portland, Oregon, Jan. 2002.
- [7] D. K. Gifford, P. Jouvelot, J. M. Lucassen, and M. A. Sheldon. FX-87 Reference Manual. Technical Report MIT/LCS/TR-407, MIT Laboratory for Computer Science, Sept. 1987.
- [8] L. Gong. *Java Security Architecture (JDK 1.2)*, 1998. <http://java.sun.com/products/jdk/1.2/docs/guide/security/spec/security-spec.doc.html>.
- [9] J. M. Lucassen and D. K. Gifford. Polymorphic Effect Systems. In *Proceedings of the 15th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 47–57, San Diego, California, Jan. 1988.
- [10] J. C. Mitchell. Abstract Types Have Existential Type. *ACM Transactions on Programming Languages and Systems*, 10(3):470–502, July 1988.
- [11] G. Necula and P. Lee. Safe, Untrusted Agents using Proof-Carrying Code. In I. G. Vigna, editor, *Special Issue on Mobile Agent Security*, volume 1419 of *Lecture Notes in Computer Science*, pages 61–91. Springer-Verlag, 1998.

$$x \in \text{Vars} \quad \ell \in \text{Locations} \quad k \in \text{KeyPairs} \cup \{\text{bot}, \text{top}\} \quad A \subseteq_{\text{finite}} \text{KeyPairs} \cup \{\text{bot}, \text{top}\}$$

raw values $r ::= m \mid \text{loc } (\ell) \mid \text{lkey } (k) \mid \text{gkey } (k) \mid \langle v_1, v_2, \dots, v_n \rangle \mid \lambda x : \bullet.e \mid \Lambda \bullet.v \mid \text{pack } v \text{ as } \bullet$
 values $v ::= x \mid k \ r \mid \text{err}$
 expressions $e ::= v \mid \langle e_1, e_2, \dots, e_n \rangle \mid e.i \mid e_1 \ e_2 \mid \text{ref } e \mid e_1 := e_2 \mid ! e \mid \text{spawn } e \mid$
 $e [\bullet] \mid \text{pack } e \text{ as } \bullet \mid \text{open } x = e_1 \text{ as } \bullet \text{ in } e_2 \mid$
 $\text{limit } e_1, e_2, \dots, e_{n-1} \text{ in } e_n \mid \text{grant } e_1 \text{ in } e_2 \mid \text{newkey} < e \mid \text{associate } e_1 \text{ with } e_2 \mid \text{access } A \text{ in } e$

Figure 5: Target language

R -contexts $R ::= [] \mid \langle \dots, v_i, R, \dots \rangle \mid R.i \mid R e \mid v R \mid \text{ref } R \mid R := e \mid v := R \mid R [\bullet] \mid \text{pack } R \text{ as } \bullet \mid \text{open } x = R \text{ as } \bullet \text{ in}$
 $\text{limit } \dots, v_i, R, \dots \text{ in } e \mid \text{grant } R \text{ in } e \mid \text{newkey} < R \mid \text{associate } R \text{ with } e \mid \text{associate } v \text{ with } R$
 E -contexts $E ::= [] \mid \langle \dots, v_i, E, \dots \rangle \mid E.i \mid E e \mid v E \mid \text{ref } E \mid E := e \mid v := E \mid E [\bullet] \mid \text{pack } E \text{ as } \bullet \mid \text{open } x = E \text{ as } \bullet \text{ in}$
 $\text{limit } \dots, v_i, E, \dots \text{ in } e \mid \text{grant } E \text{ in } e \mid \text{newkey} < E \mid \text{associate } E \text{ with } e \mid \text{associate } v \text{ with } E \mid E$
 $\text{access } A \text{ in } E$

Figure 6: Evaluation contexts

- [12] F. Pottier, C. Skalka, and S. Smith. A Systematic Approach to Static Access Control. In D. Sands, editor, *10th European Symposium on Programming*, volume 2028 of *Lecture Notes in Computer Science*, pages 30–45, Genova, Italy, 2001. Springer-Verlag.
- [13] Úlfar Erlingsson and F. B. Schneider. IRM Enforcement of Java Stack Inspection. In *Proceedings of the 2000 IEEE Symposium on Security and Privacy*, pages 246–255, Oakland, California, May 2000.
- [14] D. Walker and K. Watkins. A Type System for Expressive Security Policies. In *Proceedings of the 27th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 254–267, Boston, Massachusetts, Jan. 2000.
- [15] D. S. Wallach, A. W. Appel, and E. W. Felten. SAFKASI: A Security Mechanism for Language-based Systems. *ACM Transactions on Software Engineering and Methodology*, 9(4):341–378, Oct. 2000.

A Type Soundness

We sometimes use the symbol – to indicate “don’t care.”

A.1 Dynamic Semantics

The first step in the dynamic semantics is to translate the source program into the language shown in Figure 5. The translation involves type-erasure, replacing each name or type annotation with a \bullet . We erase names and types so that we do not need to worry about proper alpha-renaming of names during reductions. Also each program value in the source program is translated into a program value in of the target program by annotating them with bot , i.e., for each v in the source, $\text{bot } v$. Finally, the translation replaces each Top with $\text{bot lkey } (\text{top})$.

Let e be the translated target program. Then the evaluation proceeds from the initial state $(\emptyset, \emptyset, \langle e \rangle)$ followed by a sequence of small-step reductions shown in Figure 7. Each state is of the form (S, K, \vec{e}) where S is the store mapping store locations to values, K is the key-pair manager representing immediate sub-key-pair relation between key-pairs, and \vec{e} are the execution threads.

The evaluation contexts E and R are shown in Figure 6. The only difference between E and R is that R lacks $\text{access } A \text{ in } R$.

$$\begin{array}{c}
\frac{}{(S, K, \langle \dots, E[\langle v_1, v_2, \dots, v_n \rangle], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{bot } \langle v_1, v_2, \dots, v_n \rangle], \dots \rangle)} \text{[Tuple]} \\
\frac{k \in \text{Close}_K(A)}{(S, K, \langle \dots, E[\text{access } A \text{ in } R[(k \langle \dots, v_i, \dots \rangle).i]], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[v_i]], \dots \rangle)} \text{[Proj]} \\
\frac{k \in \text{Close}_K(A)}{(S, K, \langle \dots, E[\text{access } A \text{ in } R[(k \lambda x: \bullet.e) v]], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[e[v/x]]], \dots \rangle)} \text{[App]} \\
\frac{\ell \notin \text{dom}(S)}{(S, K, \langle \dots, E[\text{ref } v], \dots \rangle) \longrightarrow (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\text{bot loc } (\ell)], \dots \rangle)} \text{[Ref]} \\
\frac{k \in \text{Close}_K(A)}{(S \cup \{\ell \mapsto -\}, K, \langle \dots, E[\text{access } A \text{ in } R[(k \text{ loc } (\ell)) := v]], \dots \rangle) \longrightarrow (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\text{access } A \text{ in } R[v]], \dots \rangle)} \\
\frac{k \in \text{Close}_K(A)}{(S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\text{access } A \text{ in } R[! k \text{ loc } (\ell)]], \dots \rangle) \longrightarrow (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\text{access } A \text{ in } R[v]], \dots \rangle)} \text{[D]} \\
\frac{}{(S, K, \langle \dots, E[\text{spawn } e], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{bot } 0], \dots, \text{access } \emptyset \text{ in } e])} \text{[Spawn]} \\
\frac{k \in \text{Close}_K(A)}{(S, K, \langle \dots, E[\text{access } A \text{ in } R[(k \wedge \bullet.v) [\bullet]]], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[v]], \dots \rangle)} \text{[Inst]} \\
\frac{}{(S, K, \langle \dots, E[\text{pack } v \text{ as } \bullet], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{bot pack } v \text{ as } \bullet], \dots \rangle)} \text{[Pack]} \\
\frac{k \in \text{Close}_K(A)}{(S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{open } x = (k \text{ pack } v \text{ as } \bullet) \text{ as } \bullet \text{ in } e]], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[e[v/x]]], \dots \rangle)} \\
\frac{\{k'_1, k'_2, \dots, k'_n\} \subseteq \text{Close}_K(A)}{(S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{limit } k'_1 \text{ lkey } (k_1), k'_2 \text{ lkey } (k_2), \dots, k'_n \text{ lkey } (k_n) \text{ in } e]], \dots \rangle) \longrightarrow} \text{[Limit]} \\
\frac{}{(S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{access } \text{Close}_K(A) \cap \text{Close}_K(\{k_1, k_2, \dots, k_n\}) \text{ in } e]], \dots \rangle)} \\
\frac{k_1 \in \text{Close}_K(A)}{(S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{grant } k_1 \text{ gkey } (k_2) \text{ in } e]], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{access } A \cup \{k_2\} \text{ in } e]], \dots \rangle)} \\
\frac{k_1 \in \text{Close}_K(A) \quad k_3 \notin \text{dom}(K) \cup \{\text{bot}, \text{top}\}}{(S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{newkey} < k_1 \text{ lkey } (k_2)]]], \dots \rangle) \longrightarrow} \text{[NewKey]} \\
\frac{}{(S, K \cup \{k_3 \mapsto k_2\}, \langle \dots, E[\text{access } A \text{ in } R[\text{bot pack bot } \langle \text{bot lkey } (k_3), \text{bot gkey } (k_3) \rangle \text{ as } \bullet]], \dots \rangle)} \\
\frac{\{k_1, k_2\} \subseteq \text{Close}_K(A)}{(S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{associate } k_1 \text{ r with } k_2 \text{ lkey } (k_3)]]], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[k_3 \text{ r}], \dots \rangle)} \text{[Associate]} \\
\frac{}{(S, K, \langle \dots, E[\text{access } A \text{ in } v], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[v], \dots \rangle)} \text{[InAccess]}
\end{array}$$

Figure 7: Small-step reduction rules

$$\begin{array}{c}
\frac{\Gamma \vdash k : \rho}{\Gamma \vdash k \text{ m} : \rho \text{ int}; \emptyset} \text{ (Int'ed)} \quad \frac{\Gamma \vdash k : \rho \quad \Gamma \vdash \ell : \tau}{\Gamma \vdash k \text{ loc}(\ell) : \rho \text{ ref}(\tau); \emptyset} \text{ (Loc)} \\
\\
\frac{\Gamma \vdash k_1 : \rho_1 \quad \Gamma \vdash k_2 : \rho_2}{\Gamma \vdash k_1 \text{ lkey}(k_2) : \rho_1 \text{ lkey}(\rho_2); \emptyset} \text{ (LKey)} \quad \frac{\Gamma \vdash k_1 : \rho_1 \quad \Gamma \vdash k_2 : \rho_2}{\Gamma \vdash k_1 \text{ gkey}(k_2) : \rho_1 \text{ gkey}(\rho_2); \emptyset} \text{ (GKey)} \\
\\
\frac{\Gamma \vdash k : \rho \quad \text{For } 1 \leq i \leq n, \Gamma \vdash v_i : \tau_i; \emptyset}{\Gamma \vdash k \langle v_1, v_2, \dots, v_n \rangle : \rho \langle \tau_1, \tau_2, \dots, \tau_n \rangle; \emptyset} \text{ (Tuple'ed)} \quad \frac{\Gamma \vdash k : \rho \quad \Gamma, x : \tau_1 \vdash e : \tau_2; L \quad \Gamma, x : \tau_1 \vdash \diamond}{\Gamma \vdash k \lambda x : \bullet. e : \rho(\tau_1 \xrightarrow{L} \tau_2); \emptyset} \text{ (Fun'ed)} \\
\\
\frac{\Gamma \vdash k : \rho_3 \quad \Gamma, \rho_1 < \rho_2 \vdash v : \tau; \emptyset \quad \Gamma, \rho_1 < \rho_2 \vdash \diamond}{\Gamma \vdash k \Lambda \bullet. v : \rho_3(\forall \rho_1 < \rho_2. \tau); \emptyset} \text{ (Gen'ed)} \quad \frac{\Gamma \vdash e : \rho_1(\forall \rho_2 < \rho_3. \tau); L \quad \Gamma \vdash \rho_4 < \rho_3}{\Gamma \vdash e[\bullet] : \tau[\rho_4/\rho_2]; L \cup \{\rho_1\}} \text{ (Inst } \bullet) \\
\\
\frac{\Gamma \vdash k : \rho_4 \quad \Gamma \vdash v : \tau[\rho_3/\rho_1]; \emptyset \quad \Gamma \vdash \rho_3 < \rho_2}{\Gamma \vdash k \text{ pack } v \text{ as } \bullet : \rho_4(\exists \rho_1 < \rho_2. \tau); \emptyset} \text{ (Packed)} \\
\\
\frac{\Gamma \vdash e : \tau[\rho_3/\rho_1]; L \quad \Gamma \vdash \rho_3 < \rho_2}{\Gamma \vdash \text{pack } e \text{ as } \bullet : \perp(\exists \rho_1 < \rho_2. \tau); L} \text{ (Pack } \bullet) \quad \frac{\Gamma \vdash e_1 : \rho_1(\exists \rho_2 < \rho_3. \tau_1); L_1 \quad \Gamma, \rho_2 < \rho_3, x : \tau_1 \vdash e_2 : \tau_2; L_2 \quad \Gamma \vdash \tau_2 \quad \Gamma \vdash L_2 \quad \Gamma, \rho_2 < \rho_3, x : \tau_1 \vdash \diamond}{\Gamma \vdash \text{open } x = e_1 \text{ as } \bullet \text{ in } e_2 : \tau_2; L_1 \cup L_2 \cup \{\rho_1\}} \text{ (Open } \bullet) \\
\\
\frac{\Gamma \vdash e : \tau; L \quad \Gamma \vdash L < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}}{\Gamma \vdash \text{access } A \text{ in } e : \tau; \emptyset} \text{ (InAccess)} \\
\\
\frac{k : \rho \in \Gamma}{\Gamma \vdash k : \rho} \text{ (KeyPair } k) \quad \frac{}{\Gamma \vdash \text{top} : \top} \text{ (KeyPair top)} \quad \frac{}{\Gamma \vdash \text{bot} : \perp} \text{ (KeyPair bot)} \quad \frac{\ell : \tau \in \Gamma}{\Gamma \vdash \ell : \tau} \text{ (Location } \ell) \\
\\
\frac{\Gamma \vdash \rho \quad \Gamma \vdash \diamond}{\Gamma, k : \rho \vdash \diamond} \text{ (Env } k) \quad \frac{\Gamma \vdash \tau \quad \Gamma \vdash \diamond}{\Gamma, \ell : \tau \vdash \diamond} \text{ (Env } \ell)
\end{array}$$

Figure 8: Type checking rules: additional rules for the target language

$Close_K$ is a function such that $k \in Close_K(A)$ iff k is a sub-key-pair of some key-pair in A in the key-pair relation K . More concretely, we define the reflexive, transitive closure of K lifted with **bot** as

$$\begin{aligned}
K_{\perp}^* &= \bigcup_{n \in \omega} R^n(K \cup \{\text{bot} \mapsto k \mid k \in \text{dom}(K) \cup \{\text{top}\}\}) \\
&\quad \cup \{k \mapsto k \mid k \in \text{dom}(K) \cup \{\text{bot}, \text{top}\}\}
\end{aligned}$$

where

$$R(K) = K \cup \{k_1 \mapsto k_2 \mid k_1 \mapsto k_3 \in K \text{ and } k_3 \mapsto k_2 \in K \text{ for some } k_3\}$$

Then $Close_K$ can be defined as

$$Close_K(A) = \{k_1 \mid k_1 \mapsto k_2 \in K_{\perp}^* \text{ for some } k_2 \in A\}$$

If the state $(S, K, \langle \dots, e, \dots \rangle)$ cannot be reduced by the rules in Figure 7 but e is not a value, then e immediately reduces to a special constant **err**, i.e.

$$(S, K, \langle \dots, e, \dots \rangle) \longrightarrow (S, K, \langle \dots, \text{err}, \dots \rangle)$$

A.2 Type Soundness

The static semantics for the target language is the type checking rules from Section 3.1 plus the additional rules in Figure 8. Each type environment is now a sequence of variable-to-type bindings, names annotated

with a supertype, location-to-type bindings, and key-pair-to-name bindings.

$$\Gamma ::= \emptyset \mid \Gamma, x:\tau \mid \Gamma, \rho_1 < \rho_2 \mid \ell:\tau \mid k:\rho$$

Note that **err** has no type. The following theorem connects source programs to target programs.

Theorem 2 *Let e_1 be an expression in the source language and e_2 be the corresponding translated expression. Then if $\Gamma \vdash e_1 : \tau; L$ then $\Gamma \vdash e_2 : \tau; L$.*

Proof: By induction on the type checking derivation. \square

In the rest of the paper, we always refer to expressions in the target language unless specified otherwise. We re-state the definition of well-typed program state.

Definition 2 $\Gamma \vdash (S, K, \langle e_1, e_2, \dots, e_n \rangle)$ (read “the program state $(S, K, \langle e_1, e_2, \dots, e_n \rangle)$ is well-typed under Γ ”) iff

- (1) $\Gamma \vdash \diamond$.
- (2) For each $1 \leq i \leq n$, $\Gamma \vdash e_i : -; \emptyset$.
- (3) $\ell \in \text{dom}(S)$ iff $\ell \in \text{dom}(\Gamma)$.
- (4) $k \in \text{dom}(K)$ iff $k \in \text{dom}(\Gamma)$.
- (5) $x \notin \text{dom}(\Gamma)$.
- (6) If $\Gamma \vdash \ell : \tau$ and $S(\ell) = v$ then $\Gamma \vdash v : \tau; -$.
- (7) Suppose $\Gamma \vdash k_1 : \rho_1$ and $\Gamma \vdash k_2 : \rho_2$. Then $\Gamma \vdash \rho_1 < \rho_2$ iff $k_1 \mapsto k_2 \in K_{\perp}^*$.

Lemma 1 *If $\Gamma \vdash \diamond$ and $\Gamma \vdash e : \tau; L$ then $\Gamma \vdash \tau$ and $\Gamma \vdash L$.*

Proof: By induction on the type checking derivation. \square

Lemma 2 (Substitution) (1) *If $\Gamma, x:\tau_1 \vdash \diamond$, $\Gamma, x:\tau_1 \vdash e : \tau_2; L$, and $\Gamma \vdash v : \tau_1; \emptyset$, then $\Gamma \vdash e[v/x] : \tau_2; L$.*

(2) *If $\Gamma, \rho_1 < \rho_2 \vdash \diamond$, $\Gamma, \rho_1 < \rho_2 \vdash e : \tau; L$, and $\Gamma \vdash \rho_3 < \rho_2$, then $\Gamma \vdash e : \tau[\rho_3/\rho_1]; L[\rho_3/\rho_1]$.*

(3) *If $\Gamma, \rho_1 < \rho_2 \vdash \diamond$, $\Gamma, \rho_1 < \rho_2 \vdash \tau_1 < \tau_2$, and $\Gamma \vdash \rho_3 < \rho_2$, then $\Gamma \vdash \tau_1[\rho_3/\rho_1] < \tau_2[\rho_3/\rho_1]$.*

(4) *If $\Gamma, \rho_1 < \rho_2, x:\tau_1 \vdash \diamond$, $\Gamma, \rho_1 < \rho_2, x:\tau_1 \vdash e : \tau_2; L$, and $\Gamma \vdash \rho_3 < \rho_2$, then $\Gamma, x:\tau_1[\rho_3/\rho_1] \vdash e : \tau_2[\rho_3/\rho_1]; L[\rho_3/\rho_1]$.*

Proof: By induction on the type checking derivation. \square

Lemma 3 *Suppose $\Gamma_1 \vdash \diamond$ and $\Gamma_1 \vdash e : \tau; L$. Then for any $\Gamma_2 \supseteq \Gamma_1$ such that $\Gamma_2 \vdash \diamond$, $\Gamma_2 \vdash e : \tau; L$.*

Proof: The only non-trivial case is when Γ_2 contains more instances of name bindings, i.e., $\rho < -$'s, than Γ_1 , since type environment well-formedness conditions may fail with more names. But it is easy to see that whenever a new name is introduced in the environment, we may choose a name that has not appeared in the environment (see (Gen'ed) and (Open \bullet)). By Lemma 1, the choices do not affect the conclusions. Hence the lemma follows. \square

Lemma 4 (Replacement) *Suppose $\Gamma_1 \vdash \diamond$ and $\Gamma_1 \vdash E[e_1] : \tau; L$. Then there is a sub-derivation $\Gamma_1 \vdash e_1 : \tau_1; L_1$ for some τ_1, L_1 . Furthermore, for any e_2 and $\Gamma_2 \supseteq \Gamma_1$ such that $\Gamma_2 \vdash \diamond$ and $\Gamma_2 \vdash e_2 : \tau_1; L_1$, we have $\Gamma_2 \vdash E[e_2] : \tau; L$.*

Proof: This follows from the usual replacement argument and Lemma 3. \square

Lemma 5 *Suppose $\Gamma \vdash R[e] : \tau; L$. Let $\Gamma \vdash e : \tau_1; L_1$ be a sub-derivation. Then $\Gamma \vdash L_1 < L$.*

Proof: By induction on the type checking derivation. \square

Lemma 6 *If $\Gamma \vdash v : \tau; L$ then $\Gamma \vdash v : \tau; \emptyset$.*

Proof: By induction on the type checking derivation. \square

We are now prepared prove our main theorem, re-stated here.

Theorem 3 (Subject Reduction) *If $\Gamma \vdash (S, K, \vec{e})$ and $(S, K, \vec{e}) \longrightarrow (S', K', \vec{e}')$, then there is Γ' such that $\Gamma' \vdash (S', K', \vec{e}')$.*

Proof: Firstly, we show that for each $e_i \in \vec{e}$, if e_i is not a value then (S, K, \vec{e}) is of the form matching the left hand side of \longrightarrow in one of the reduction rules from Figure 7 with e_i as the evaluating thread. To see this, first note that for [Proj], [App], [Ref], [Assign], [Deref], [Open], [Limit], [Grant], [NewKey], and [Associate], the redex indeed does appear in some context of the form $E[\text{access } A \text{ in } R[\]]$ since otherwise it must appear in some context of the form $R[\]$ but then by Lemma 5 $\Gamma \not\vdash e_i : -; \emptyset$. Secondly, for [Assign] (ref. [Deref]), the location assigned (ref. dereferenced) must be in the store because otherwise it violates Definition 2 (3).

Therefore it suffices to show the statement in the theorem holds for each forms matching (S, K, \vec{e}) .

[Tuple] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\langle v_1, v_2, \dots, v_n \rangle], \dots \rangle)$$

So we have $\Gamma \vdash E[\langle v_1, v_2, \dots, v_n \rangle] : -; \emptyset$. By inspection of the type checking rules, there must be a sub-derivation with (Tuple) as the last rule.

$$\frac{\text{For } 1 \leq i \leq n, \Gamma \vdash v_i : \tau_i; L_i}{\Gamma \vdash \langle v_1, v_2, \dots, v_n \rangle : \perp \langle \tau_1, \tau_2, \dots, \tau_n \rangle; \bigcup_{1 \leq i \leq n} L_i}$$

We have

$$(S, K, \langle \dots, E[\langle v_1, v_2, \dots, v_n \rangle], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{bot } \langle v_1, v_2, \dots, v_n \rangle], \dots \rangle)$$

Let $\Gamma' = \Gamma$. We have by (Tuple'ed)

$$\frac{\Gamma' \vdash \text{bot} : \perp \quad \text{For } 1 \leq i \leq n, \Gamma' \vdash v_i : \tau_i; \emptyset}{\Gamma' \vdash \text{bot } \langle v_1, v_2, \dots, v_n \rangle : \perp \langle \tau_1, \tau_2, \dots, \tau_n \rangle; \emptyset}$$

Then by Lemma 1 and (Subsumption),

$$\Gamma' \vdash \text{bot } \langle v_1, v_2, \dots, v_n \rangle : \perp \langle \tau_1, \tau_2, \dots, \tau_n \rangle; \bigcup_{1 \leq i \leq n} L_i$$

Hence by Lemma 4, it follows that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{bot } \langle v_1, v_2, \dots, v_n \rangle], \dots \rangle)$$

[Proj] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[(k \langle \dots, v_i, \dots \rangle).i]], \dots \rangle)$$

So we have $\Gamma \vdash E[\text{access } A \text{ in } R[(k \langle \dots, v_i, \dots \rangle).i]] : -; \emptyset$. By inspection of the type checking rules, there must be a sub-derivation with (Proj) as the last rule.

$$\frac{\Gamma \vdash k \langle \dots, v_i, \dots \rangle : \rho \langle \tau_1, \tau_2, \dots, \tau_n \rangle; L}{\Gamma \vdash (k \langle \dots, v_i, \dots \rangle).i : \tau_i; L \cup \{\rho\}}$$

We first show that

$$(S, K, \langle \dots, E[\text{access } A \text{ in } R[(k \langle \dots, v_i, \dots \rangle).i]], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[v_i]], \dots \rangle)$$

It suffices to show that $k \in \text{Close}_K(A)$. By Lemma 5 and (InAccess), it must be the case that

$$\Gamma \vdash \{\rho\} < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}$$

By inspection of the type checking rules, we see that there must be ρ' such that

$$\Gamma \vdash k : \rho' \quad \Gamma \vdash \rho' < \rho$$

So by Definition 2 (7), $k \in \text{Close}_K(A)$.

Now let $\Gamma' = \Gamma$. We next show that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[v_i]], \dots \rangle)$$

By inspection of the type checking rules, we see that the sub-derivation of the premise must end with (Tuple'ed) followed by zero or more (Subsumption)'s. So

$$\Gamma' \vdash v_i : \tau'_i; \emptyset \quad \Gamma' \vdash \tau'_i < \tau_i \quad \Gamma' \vdash \emptyset < L \cup \{\rho\}$$

So by Lemma 1 and (Subsumption)

$$\Gamma' \vdash v_i : \tau_i; L \cup \{\rho\}$$

So by Lemma 4 it follows that,

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[v_i]], \dots \rangle)$$

[App] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[(k \lambda x : \bullet.e) v]], \dots \rangle)$$

So we have $\Gamma \vdash E[\text{access } A \text{ in } R[(k \lambda x : \bullet.e) v]] : -; \emptyset$. By inspection of the type checking rules, there must be a sub-derivation with (App) as the last rule.

$$\frac{\Gamma \vdash k \lambda x : \bullet.e : \rho \ (\tau_1 \xrightarrow{L_1} \tau_2); L_2 \quad \Gamma \vdash v : \tau_1; L_3}{\Gamma \vdash k \lambda x : \bullet.e v : \tau_2; L_1 \cup L_2 \cup L_3 \cup \{\rho\}}$$

We first show that

$$(S, K, \langle \dots, E[\text{access } A \text{ in } R[(k \lambda x : \bullet.e) v]], \dots \rangle \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[e[v/x]]], \dots \rangle)$$

It suffices to show that $k \in \text{Close}_K(A)$. By Lemma 5 and (InAccess), it must be the case that

$$\Gamma \vdash \{\rho\} < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}$$

By inspection of the type checking rules, we see that there must be ρ' such that

$$\Gamma \vdash k : \rho' \quad \Gamma \vdash \rho' < \rho$$

So by Definition 2 (7), $k \in \text{Close}_K(A)$.

Now let $\Gamma' = \Gamma$. We next show that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[e[v/x]]], \dots \rangle)$$

By inspection of the type checking rules, we see that the sub-derivation of the first premise must end with (Fun'ed) followed by zero or more (Subsumption)'s. So

$$\begin{array}{l} \Gamma', x : \tau'_1 \vdash e : \tau'_2; L'_1 \quad \Gamma', x : \tau'_1 \vdash \diamond \\ \Gamma' \vdash \tau_1 < \tau'_1 \quad \Gamma' \vdash \tau'_2 < \tau'_2 \\ \Gamma' \vdash L'_1 < L_1 \quad \Gamma' \vdash \emptyset < L_2 \end{array}$$

So by (Subsumption),

$$\Gamma' \vdash v : \tau'_1; L_3$$

and by Lemma 2 (1), Lemma 1, and (Subsumption),

$$\Gamma' \vdash e[v/x] : \tau_2; L_1 \cup L_2 \cup L_3$$

So by Lemma 4 it follows that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[e[v/x]]], \dots \rangle)$$

[Ref] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\mathbf{ref} \ v], \dots \rangle)$$

So we have $\Gamma \vdash E[\mathbf{ref} \ v] : -; \emptyset$. By inspection of the type checking rules, there must be a sub-derivation with (Ref) as the last rule.

$$\frac{\Gamma \vdash v : \tau; L}{\Gamma \vdash \mathbf{ref} \ v : \perp \ \mathit{ref}(\tau); L}$$

We have

$$(S, K, \langle \dots, E[\mathbf{ref} \ v], \dots \rangle) \longrightarrow (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\mathbf{bot} \ \mathit{loc}(\ell)], \dots \rangle)$$

for $\ell \notin \mathit{dom}(S)$. Let $\Gamma' = \Gamma, \ell : \tau$. We have by (Loc)

$$\frac{\Gamma' \vdash \mathbf{bot} : \perp \quad \Gamma' \vdash \ell : \tau}{\Gamma' \vdash \mathbf{bot} \ \mathit{loc}(\ell) : \perp \ \mathit{ref}(\tau); \emptyset}$$

Then by Lemma 1 and (Subsumption),

$$\Gamma' \vdash \mathbf{bot} \ \mathit{loc}(\ell) : \perp \ \mathit{ref}(\tau); L$$

Now, we have $\Gamma' \vdash \ell : \tau$, $(S \cup \{\ell \mapsto v\})(\ell) = v$, and $\Gamma \vdash v : \tau; -$. Hence by Lemma 3 and Lemma 4, it follows that

$$\Gamma' \vdash (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\mathbf{bot} \ \mathit{loc}(\ell)], \dots \rangle)$$

[Assign] Suppose

$$\Gamma \vdash (S \cup \{\ell \mapsto -\}, K, \langle \dots, E[\mathbf{access} \ A \ \mathbf{in} \ R[(k \ \mathit{loc}(\ell)) := v]], \dots \rangle)$$

So we have $\Gamma \vdash E[\mathbf{access} \ A \ \mathbf{in} \ R[(k \ \mathit{loc}(\ell)) := v]] : -; \emptyset$. By inspection of the type checking rules, there must be a sub-derivation with (Assign) as the last rule.

$$\frac{\Gamma \vdash k \ \mathit{loc}(\ell) : \rho \ \mathit{ref}(\tau); L_1 \quad \Gamma \vdash v : \tau; L_2}{\Gamma \vdash (k \ \mathit{loc}(\ell)) := v : \tau; L_1 \cup L_2 \cup \{\rho\}}$$

We first show that

$$(S \cup \{\ell \mapsto -\}, K, \langle \dots, E[\mathbf{access} \ A \ \mathbf{in} \ R[(k \ \mathit{loc}(\ell)) := v]], \dots \rangle) \longrightarrow (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\mathbf{access} \ A \ \mathbf{in} \ R[v]], \dots \rangle)$$

It suffices to show that $k \in \mathit{Close}_K(A)$. By Lemma 5 and (InAccess), it must be the case that

$$\Gamma \vdash \{\rho\} < \{\rho \mid \Gamma \vdash k : \rho \ \text{for some } k \in A\}$$

By inspection of the type checking rules, we see that there must be ρ' such that

$$\Gamma \vdash k : \rho' \quad \Gamma \vdash \rho' < \rho$$

So by Definition 2 (7), $k \in \mathit{Close}_K(A)$.

Now let $\Gamma' = \Gamma$. We next show that

$$\Gamma' \vdash (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\mathbf{access} \ A \ \mathbf{in} \ R[v]], \dots \rangle)$$

By inspection of the type checking rules, we see that the sub-derivation of the first premise must end with (Loc) followed by zero or more (Subsumption)'s. So

$$\Gamma' \vdash \ell : \tau$$

(Subtyping is invariant under pointer types.) Hence Definition 2 (6) is satisfied. Also by Lemma 1 and (Subsumption)

$$\Gamma \vdash v : \tau; L_1 \cup L_2 \cup \{\rho\}$$

Hence by Lemma 4, it follows that

$$\Gamma' \vdash (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\mathbf{access} \ A \ \mathbf{in} \ R[v]], \dots \rangle)$$

[Deref] Suppose

$$\Gamma \vdash (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\text{access } A \text{ in } R[! k \text{ loc } (\ell)]], \dots \rangle)$$

So we have $\Gamma \vdash E[\text{access } A \text{ in } R[! k \text{ loc } (\ell)]] : -; \emptyset$. By inspection of the type checking rules, there must be a sub-derivation with (Deref) as the last rule.

$$\frac{\Gamma \vdash k \text{ loc } (\ell) : \rho \text{ ref } (\tau); L}{\Gamma \vdash ! k \text{ loc } (\ell) : \tau; L \cup \{\rho\}}$$

We first show that

$$\begin{aligned} & (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\text{access } A \text{ in } R[! k \text{ loc } (\ell)]], \dots \rangle) \\ & \longrightarrow (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\text{access } A \text{ in } R[v]], \dots \rangle) \end{aligned}$$

It suffices to show that $k \in \text{Close}_K(A)$. By Lemma 5 and (InAccess), it must be the case that

$$\Gamma \vdash \{\rho\} < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}$$

By inspection of the type checking rules, we see that there must be ρ' such that

$$\Gamma \vdash k : \rho' \quad \Gamma \vdash \rho' < \rho$$

So by Definition 2 (7), $k \in \text{Close}_K(A)$.

Now let $\Gamma' = \Gamma$. We next show that

$$\Gamma' \vdash (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\text{access } A \text{ in } R[v]], \dots \rangle)$$

By inspection of the type checking rules, we see that the sub-derivation of the premise must end with (Loc) followed by zero or more (Subsumption)'s. So

$$\Gamma' \vdash \ell : \tau$$

(Subtyping is invariant under pointer types.) Hence by Definition 2 (6) and Lemma 6, it must be the case that

$$\Gamma \vdash v : \tau; \emptyset$$

So by Lemma 1 and (Subsumption)

$$\Gamma \vdash v : \tau; L \cup \{\rho\}$$

Hence by Lemma 4, it follows that

$$\Gamma' \vdash (S \cup \{\ell \mapsto v\}, K, \langle \dots, E[\text{access } A \text{ in } R[v]], \dots \rangle)$$

[Spawn] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\text{spawn } e], \dots \rangle)$$

So we have $\Gamma \vdash E[\text{spawn } e] : -; \emptyset$. By inspection of the type checking rules, there must be a sub-derivation with (Spawn) as the last rule.

$$\frac{\Gamma \vdash e : \tau; \emptyset}{\Gamma \vdash \text{spawn } e : \perp \text{ int}; \emptyset}$$

We have

$$\begin{aligned} & (S, K, \langle \dots, E[\text{spawn } e], \dots \rangle) \longrightarrow \\ & (S, K, \langle \dots, E[\text{bot } 0], \dots, \text{access } \emptyset \text{ in } e \rangle) \end{aligned}$$

Let $\Gamma' = \Gamma$. We have by (Int'ed)

$$\frac{\Gamma' \vdash \text{bot} : \perp}{\Gamma' \vdash \text{bot } 0 : \perp \text{ int}; \emptyset}$$

Also by (InAccess),

$$\frac{\Gamma' \vdash e : \tau; \emptyset \quad \Gamma' \vdash \emptyset < \emptyset}{\Gamma' \vdash \text{access } \emptyset \text{ in } e : \tau; \emptyset}$$

Hence by Lemma 4, it follows that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{bot } 0], \dots, \text{access } \emptyset \text{ in } e \rangle)$$

[Inst] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[(k \Lambda \bullet .v) [\bullet]], \dots] \rangle)$$

So we have $\Gamma \vdash E[\text{access } A \text{ in } R[(k \Lambda \bullet .v) [\bullet]]] : -; \emptyset$. By inspection of the type checking rules, there must be a sub-derivation with (Inst \bullet) as the last rule.

$$\frac{\Gamma \vdash k \Lambda \bullet .v : \rho_1 (\forall \rho_2 < \rho_3. \tau); L \quad \Gamma \vdash \rho_4 < \rho_3}{\Gamma \vdash (k \Lambda \bullet .v) [\bullet] : \tau[\rho_4/\rho_2]; L \cup \{\rho_1\}}$$

We first show that

$$(S, K, \langle \dots, E[\text{access } A \text{ in } R[(k \Lambda \bullet .v) [\bullet]], \dots] \rangle \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[v]], \dots \rangle)$$

It suffices to show that $k \in \text{Close}_K(A)$. By Lemma 5 and (InAccess), it must be the case that

$$\Gamma \vdash \{\rho\} < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}$$

By inspection of the type checking rules, we see that there must be ρ' such that

$$\Gamma \vdash k : \rho' \quad \Gamma \vdash \rho' < \rho$$

So by Definition 2 (7), $k \in \text{Close}_K(A)$.

Now let $\Gamma' = \Gamma$. We next show that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[v]], \dots \rangle)$$

By inspection of the type checking rules, we see that the sub-derivation of the first premise must end with (Gen'ed) followed by zero or more (Subsumption)'s. So

$$\begin{array}{c} \Gamma', \rho_2 < \rho'_3 \vdash v : \tau'; \emptyset \quad \Gamma', \rho_2 < \rho'_3 \vdash \diamond \\ \Gamma' \vdash \rho_3 < \rho'_3 \quad \Gamma' \vdash \tau' < \tau \\ \Gamma' \vdash \emptyset < L \end{array}$$

So we have

$$\Gamma' \vdash \rho_4 < \rho'_3$$

and by Lemma 2 (2), Lemma 1, and (Subsumption),

$$\Gamma' \vdash v : \tau[\rho_4/\rho_2]; L \cup \{\rho_1\}$$

So by Lemma 4 it follows that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[v]], \dots \rangle)$$

[Pack] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\text{pack } v \text{ as } \bullet], \dots \rangle)$$

So we have $\Gamma \vdash E[\text{pack } v \text{ as } \bullet] : -; \emptyset$. By inspection of the type checking rules, there must be a sub-derivation with (Pack \bullet) as the last rule.

$$\frac{\Gamma \vdash v : \tau[\rho_3/\rho_1]; L \quad \Gamma \vdash \rho_3 < \rho_2}{\Gamma \vdash \text{pack } v \text{ as } \bullet : \perp (\exists \rho_1 < \rho_2. \tau); L}$$

We have

$$(S, K, \langle \dots, E[\text{pack } v \text{ as } \bullet], \dots \rangle \longrightarrow (S, K, \langle \dots, E[\text{bot pack } v \text{ as } \bullet], \dots \rangle)$$

Let $\Gamma' = \Gamma$. We have by (Packed)

$$\frac{\Gamma' \vdash \text{bot} : \perp \quad \Gamma' \vdash v : \tau[\rho_3/\rho_1]; \emptyset \quad \Gamma' \vdash \rho_3 < \rho_2}{\Gamma' \vdash \text{bot pack } v \text{ as } \bullet : \perp (\exists \rho_1 < \rho_2. \tau); \emptyset}$$

Then by Lemma 1 and (Subsumption),

$$\Gamma' \vdash \text{bot pack } v \text{ as } \bullet : \perp (\exists \rho_1 < \rho_2. \tau); L$$

Hence by Lemma 4, it follows that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{bot pack } v \text{ as } \bullet], \dots \rangle)$$

[Open] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{open } x = (k \text{ pack } v \text{ as } \bullet) \text{ as } \bullet \text{ in } e]], \dots \rangle)$$

So we have

$$\Gamma \vdash E[\text{access } A \text{ in } R[\text{open } x = (k \text{ pack } v \text{ as } \bullet) \text{ as } \bullet \text{ in } e]] : -; \emptyset$$

By inspection of the type checking rules, there must be a sub-derivation with (Open \bullet) as the last rule.

$$\frac{\Gamma \vdash (k \text{ pack } v \text{ as } \bullet) : \rho_1 \ (\exists \rho_2 < \rho_3. \tau_1); L_1 \quad \Gamma, \rho_2 < \rho_3, x : \tau_1 \vdash e : \tau_2; L_2 \quad \Gamma \vdash \tau_2 \quad \Gamma \vdash L_2 \quad \Gamma, \rho_2 < \rho_3, x : \tau_1 \vdash \diamond}{\Gamma \vdash \text{open } x = (k \text{ pack } v \text{ as } \bullet) \text{ as } \bullet \text{ in } e : \tau_2; L_1 \cup L_2 \cup \{\rho_1\}}$$

We first show that

$$\begin{aligned} & (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{open } x = (k \text{ pack } v \text{ as } \bullet) \text{ as } \bullet \text{ in } e]], \dots \rangle) \\ & \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[e[v/x]]], \dots \rangle) \end{aligned}$$

It suffices to show that $k \in \text{Close}_K(A)$. By Lemma 5 and (InAccess), it must be the case that

$$\Gamma \vdash \{\rho_1\} < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}$$

By inspection of the type checking rules, we see that there must be ρ'_1 such that

$$\Gamma \vdash k : \rho'_1 \quad \Gamma \vdash \rho'_1 < \rho_1$$

So by Definition 2 (7), $k \in \text{Close}_K(A)$.

Now let $\Gamma' = \Gamma$. We next show that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[e[v/x]]], \dots \rangle)$$

By inspection of the type checking rules, we see that the sub-derivation of the first premise must end with (Packed) followed by zero or more (Subsumption)'s. That is, for some ρ_4

$$\begin{aligned} & \Gamma' \vdash v : \tau'_1[\rho_4/\rho_2]; \emptyset \\ & \Gamma' \vdash \rho_4 < \rho'_3 \quad \Gamma' \vdash \rho'_3 < \rho_3 \\ & \Gamma', \rho_2 < \rho'_3 \vdash \tau'_1 < \tau_1 \end{aligned}$$

So we have by Lemma 2 (3)

$$\Gamma' \vdash \tau'_1[\rho_4/\rho_2] < \tau_1[\rho_4/\rho_2]$$

so

$$\Gamma' \vdash v : \tau_1[\rho_4/\rho_2]; \emptyset$$

Furthermore, because $\Gamma' \vdash \tau_2$, $\Gamma' \vdash L_2$ and $\Gamma' \vdash \rho_4 < \rho_2$, we have by Lemma 2 (4)

$$\Gamma', x : \tau_1[\rho_4/\rho_2] \vdash e : \tau_2; L_2$$

Hence by Lemma 2 (1)

$$\Gamma' \vdash e[v/x] : \tau_2; L_2$$

Finally by Lemma 4 it follows that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[e[v/x]]], \dots \rangle)$$

[Limit] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{limit } k'_1 \text{ lkey } (k_1), \dots, k'_n \text{ lkey } (k_n) \text{ in } e]], \dots \rangle)$$

So we have

$$\Gamma \vdash E[\text{access } A \text{ in } R[\text{limit } k'_1 \text{ lkey } (k_1), \dots, k'_n \text{ lkey } (k_n) \text{ in } e]] : -; \emptyset$$

By inspection of the type checking rules, there must be a sub-derivation with (Limit) as the last rule.

$$\frac{\text{For } 1 \leq i \leq n, \Gamma \vdash k'_i \text{ lkey } (k_i) : \rho'_i \text{ lkey } (\rho_i); L_i \quad \Gamma \vdash e : \tau; L \quad \Gamma \vdash L < \{\rho_1, \rho_2, \dots, \rho_n\}}{\Gamma \vdash \text{limit } k'_1 \text{ lkey } (k_1), \dots, k'_n \text{ lkey } (k_n) \text{ in } e : \tau; L \cup \bigcup_{1 \leq i \leq n-1} (L_i \cup \{\rho'_i\})}$$

Let $B = \text{Close}_K(A) \cap \text{Close}_K(\{k_1, k_2, \dots, k_n\})$. We first show that

$$(S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{limit } k'_1 \text{ lkey } (k_1), \dots, k'_n \text{ lkey } (k_n) \text{ in } e]], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{access } B \text{ in } e]], \dots \rangle)$$

It suffices to show that for $1 \leq i \leq n$, $k'_i \in \text{Close}_K(A)$. By Lemma 5 and (InAccess), it must be the case that

$$\Gamma \vdash \{\rho'_i\} < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}$$

By inspection of the type checking rules, we see that there must be ρ''_i such that

$$\Gamma \vdash k'_i : \rho''_i \quad \Gamma \vdash \rho''_i < \rho'_i$$

So by Definition 2 (7), $k'_i \in \text{Close}_K(A)$.

Now let $\Gamma' = \Gamma$. We next show that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{access } B \text{ in } e]], \dots \rangle)$$

Firstly, by Lemma 5

$$\Gamma' \vdash L < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}$$

Hence by $\Gamma' \vdash L < \{\rho_1, \rho_2, \dots, \rho_n\}$ and Definition 2 (7), it follows that

$$\Gamma' \vdash L < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in B\}$$

(Since subtyping is invariant for limit key types.) Then by Lemma 1 and (Subsumption),

$$\Gamma' \vdash \text{access } B \text{ in } e : \tau; L \cup \bigcup_{1 \leq i \leq n-1} (L_i \cup \{\rho'_i\})$$

Hence by Lemma 4, it follows that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{access } B \text{ in } e]], \dots \rangle)$$

[Grant] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{grant } k_1 \text{ gkey } (k_2) \text{ in } e]], \dots \rangle)$$

So we have $\Gamma \vdash E[\text{access } A \text{ in } R[\text{grant } k_1 \text{ gkey } (k_2) \text{ in } e]] : -; \emptyset$. By inspection of the type checking rules, there must be a sub-derivation with (Grant) as the last rule.

$$\frac{\Gamma \vdash k_1 \text{ gkey } (k_2) : \rho_1 \text{ gkey } (\rho_2); L_1 \quad \Gamma \vdash e : \tau; L_2 \cup \{\rho_2\}}{\Gamma \vdash \text{grant } k_1 \text{ gkey } (k_2) \text{ in } e : \tau; L_1 \cup L_2 \cup \{\rho_1\}}$$

We first show that

$$(S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{grant } k_1 \text{ gkey } (k_2) \text{ in } e]], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{access } A \cup \{k_2\} \text{ in } e]], \dots \rangle)$$

It suffices to show that $k_1 \in \text{Close}_K(A)$. By Lemma 5 and (InAccess), it must be the case that

$$\Gamma \vdash \{\rho_1\} < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}$$

By inspection of the type checking rules, we see that there must be ρ' such that

$$\Gamma \vdash k_1 : \rho' \quad \Gamma \vdash \rho' < \rho_1$$

So by Definition 2 (7), $k_1 \in \text{Close}_K(A)$.

Now let $\Gamma' = \Gamma'$. We next show that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{access } A \cup \{k_2\} \text{ in } e]], \dots \rangle)$$

By Lemma 5

$$\Gamma' \vdash L < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}$$

Therefore by Definition 2 (7), it follows that

$$\Gamma' \vdash L \cup \{\rho_2\} < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A \cup \{k_2\}\}$$

(Since subtyping is invariant for grant key types.) Then by Lemma 1 and (Subsumption),

$$\Gamma' \vdash \text{access } A \cup \{k_2\} \text{ in } e : \tau; L_1 \cup L_2 \cup \{\rho_2\}$$

Hence by Lemma 4, it follows that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{access } A \cup \{k_2\} \text{ in } e]], \dots \rangle)$$

[NewKey] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{newkey} < k_1 \text{ lkey } (k_2)]], \dots \rangle)$$

So we have $\Gamma \vdash E[\text{access } A \text{ in } R[\text{newkey} < k_1 \text{ lkey } (k_2)]] : -; \emptyset$. By inspection of the type checking rules, there must be a sub-derivation with (NewKey) as the last rule.

$$\frac{\Gamma \vdash k_1 \text{ lkey } (k_2) : \rho_1 \text{ lkey } (\rho_2); L}{\Gamma \vdash \text{newkey} < k_1 \text{ lkey } (k_2) : \perp (\exists \rho_3 < \rho_2. \perp (\perp \text{lkey } (\rho_3), \perp \text{gkey } (\rho_3))); L \cup \{\rho_1\}}$$

We first show that

$$(S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{newkey} < k_1 \text{ lkey } (k_2)]], \dots \rangle) \longrightarrow (S, K \cup \{k_3 \mapsto k_2\}, \langle \dots, E[\text{access } A \text{ in } R[\text{bot pack bot } (\text{bot lkey } (k_3), \text{bot gkey } (k_3)) \text{ as } \bullet]], \dots \rangle)$$

for some $k_3 \notin \text{dom}(K) \cup \{\text{bot}, \text{top}\}$. It suffices to show that $k_1 \in \text{Close}_K(A)$. By Lemma 5 and (InAccess), it must be the case that

$$\Gamma \vdash \{\rho_1\} < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}$$

By inspection of the type checking rules, we see that there must be ρ' such that

$$\Gamma \vdash k_1 : \rho' \quad \Gamma \vdash \rho' < \rho_1$$

So by Definition 2 (7), $k_1 \in \text{Close}_K(A)$.

Pick $\rho_3 \notin \text{dom}(\Gamma) \cup \{\perp, \top\}$. Let $\Gamma' = \Gamma, \rho_3 < \rho_2, k : \rho_3$. We next show that

$$\Gamma' \vdash (S, K \cup \{k_3 \mapsto k_2\}, \langle \dots, E[\text{access } A \text{ in } R[\text{bot pack bot } \langle \text{bot lkey } (k_3), \text{bot gkey } (k_3) \rangle \text{ as } \bullet], \dots] \rangle)$$

By (Packed) and (Tuple'ed), we have

$$\Gamma' \vdash \text{bot pack bot } \langle \text{bot lkey } (k_3), \text{bot gkey } (k_3) \rangle \text{ as } \bullet : \perp (\exists \rho_3 < \rho_2. \perp (\perp \text{lkey } (\rho_2), \perp \text{gkey } (\rho_2))); \emptyset$$

Hence by Lemma 1 and (Subsumption),

$$\Gamma' \vdash \text{bot pack bot } \langle \text{bot lkey } (k_3), \text{bot gkey } (k_3) \rangle \text{ as } \bullet : \perp (\exists \rho_3 < \rho_2. \perp (\perp \text{lkey } (\rho_2), \perp \text{gkey } (\rho_2))); L \cup \{\rho_1\}$$

Also, it is easy to see that Definition 2 (4) and (7) is satisfied. Hence

$$\Gamma' \vdash (S, K \cup \{k_3 \mapsto k_2\}, \langle \dots, E[\text{access } A \text{ in } R[\text{bot pack bot } \langle \text{bot lkey } (k_3), \text{bot gkey } (k_3) \rangle \text{ as } \bullet], \dots] \rangle)$$

[Associate] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{associate } k_1 r \text{ with } k_2 \text{lkey } (k_3)], \dots] \rangle)$$

So we have

$$\Gamma \vdash E[\text{access } A \text{ in } R[\text{associate } k_1 r \text{ with } k_2 \text{lkey } (k_3)]] : -; \emptyset$$

By inspection of the type checking rules, there must be a sub-derivation with (Associate) as the last rule.

$$\frac{\Gamma \vdash k_1 r : \rho_1 \sigma; L_1 \quad \Gamma \vdash k_2 \text{lkey } (k_3) : \rho_2 \text{lkey } (\rho_3); L_2}{\Gamma \vdash \text{associate } k_1 r \text{ with } k_2 \text{lkey } (k_3) : \rho_3 \sigma; L_1 \cup L_2 \cup \{\rho_1, \rho_2\}}$$

We first show that

$$(S, K, \langle \dots, E[\text{access } A \text{ in } R[\text{associate } k_1 r \text{ with } k_2 \text{lkey } (k_3)], \dots] \rangle) \longrightarrow (S, K, \langle \dots, E[\text{access } A \text{ in } R[k_3 r], \dots] \rangle)$$

It suffices to show that $k_1 \in \text{Close}_K(A)$ and $k_2 \in \text{Close}_K(A)$. By Lemma 5 and (InAccess), it must be the case that

$$\Gamma \vdash \{\rho_1\} < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}$$

By inspection of the type checking rules, we see that there must be ρ' such that

$$\Gamma \vdash k_1 : \rho' \quad \Gamma \vdash \rho' < \rho_1$$

So by Definition 2 (7), $k_1 \in \text{Close}_K(A)$. The case for k_2 is analogous.

Now let $\Gamma' = \Gamma$. We next show that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[k_3 r], \dots] \rangle)$$

The cases split between (Int'ed), (Loc), (Lkey), (Gkey), (Tuple'ed), (Fun'ed), (Gen'ed), and (Packed) depending on σ , but in each case we have

$$\Gamma' \vdash k_3 r : \rho_3 \sigma; \emptyset$$

(Since subtyping is invariant for limit key types.) Then by Lemma 1 and (Subsumption),

$$\Gamma' \vdash k_3 r : \rho_3 \sigma; L_1 \cup L_2 \cup \{\rho_1, \rho_2\}$$

Hence by Lemma 4, it follows that

$$\Gamma' \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } R[k_3 r], \dots] \rangle)$$

[InAccess] Suppose

$$\Gamma \vdash (S, K, \langle \dots, E[\text{access } A \text{ in } v], \dots \rangle)$$

So we have $\Gamma \vdash E[\text{access } A \text{ in } v] : -; \emptyset$. By inspection of the type checking rules, there must be a sub-derivation with (InAccess) as the last rule.

$$\frac{\Gamma \vdash v : \tau; L \quad \Gamma \vdash L < \{\rho \mid \Gamma \vdash k : \rho \text{ for some } k \in A\}}{\Gamma \vdash \text{access } A \text{ in } v : \tau; \emptyset}$$

We have

$$(S, K, \langle \dots, E[\text{access } A \text{ in } v], \dots \rangle) \longrightarrow (S, K, \langle \dots, E[v], \dots \rangle)$$

Let $\Gamma' = \Gamma$. Then by Lemma 6,

$$\Gamma \vdash v : \tau; \emptyset$$

Hence by Lemma 4, it follows that

$$\Gamma' \vdash (S, K, \langle \dots, E[v], \dots \rangle)$$

This concludes the proof of Theorem 3. □

Let e be a source program. We say that e gets stuck if it cannot be translated to the target language or if its evaluation reaches a state of the form $(S, K, \langle \dots, err, \dots \rangle)$.

Corollary 1 (Type Soundness) *If the program e is well-typed then it does not get stuck.*

Proof: Recall that each source program e is well-typed iff $\emptyset \vdash e : -; \emptyset$. The corollary follows from Theorem 2 and Theorem 3 since the initial state $(\emptyset, \emptyset, \langle e \rangle)$ is well-typed under the type environment \emptyset . □ The corollary implies that a well-typed program does not violate its access control policy.