is this information that a game client would typically send today and that we permit the client to omit in our approach.

Due to the scope of what it tries to detect, however, our technique has some limitations that are immediately evident. First, our technique will not detect cheats that are permitted by the sanctioned client software due to bugs. Second, modifications to the game client that do not change its behavior as seen at the server will go unnoticed by our technique. For example, any action that is possible to perform will be accepted, and so cheating by modifying the client program to make difficult (but possible) actions easy will go undetected. Put in a more positive light, however, this means that our technique has no false alarms, assuming that symbolic execution successfully explores all paths through the client. As another example, a client modification that discloses information to the player that should be hidden, e.g., such as a common cheat that uncovers parts of the game map that should be obscured, will go unnoticed by our technique. In the limit, a player could write his own version of the game client from scratch and still go undetected, provided that the behaviors it emits, as witnessed by the server, are a subset of those that the sanctioned client software could emit.

4. OUR APPROACH

Our detection mechanism analyzes client output (as seen by the game server) and determines whether that output could in fact have been produced by a valid game client. Toward that end, a key step of our approach is to profile the game client's source code using symbolic execution and then use the results in our analysis of observed client outputs. We begin with a summary of symbolic execution in Section 4.1, and then discuss its application in our context in Sections 4.2–4.6. The symbolic execution engine that we use in our work is KLEE [Cadar et al. 2008], with some modifications to make it more suitable for our task.

Before we continue, we clarify our use of certain terminology. Below, when we refer to a *valid* client, we mean a client that faithfully executes a sanctioned game-client program (and does not interfere with its behavior). Values or messages are then valid if they could have been emitted by a valid game client.

4.1. Symbolic Execution

Symbolic execution is a way of "executing" a program while exploring all execution paths, for example, to find bugs in the program. Symbolic execution works by executing the software with its initial inputs specially marked so they are allowed to be "anything"; the memory regions of the input are marked as symbolic and are not given any initial value. The program is executed step-by-step, building constraints on the symbolic variables based on the program's operations on those variables. For example, if the program sets $a \leftarrow b + c$, where a, b, and c are all marked as symbolic, then after the operation, there will be a new logical constraint on the value of a that states that it must equal the sum of b and c. When the program conditionally branches on a symbolic value, execution forks and both program branches are followed, with the true branch forming a constraint that the symbolic value evaluates to true and the false branch forming the opposite constraint. Using this strategy, symbolic execution attempts to follow each possible code path in the target program, building a constraint that must hold on execution of that path.

Symbolic execution can help locate software bugs by providing constraints that enable a constraint solver (KLEE uses STP [Ganesh and Dill 2007]) to generate concrete inputs that cause errors to occur. For example, if execution reaches an error condition (or a state thought to be "impossible"), then a constraint solver can use the constraints associated with that path to solve for a concrete input value which triggers the error

```
100: loc ← 0;
101:
102: while true do
103:    key ← readkey();
104:    if key = ESC then
105:       endgame();
106:    else if key = '↑' then
107:       loc ← loc + 1;
108:    else if key = '↓' then
109:       loc ← loc − 1;
110:    end if
111:    sendlocation(loc);
112: end while
```

(a) A toy game client ...

```
200: prev_loc ← symbolic;
201: loc ← prev_loc;
202: while true do
203:    key ← symbolic;
204:    if key = ESC then
205:       endgame();
206:    else if key = '↑' then
207:       loc ← loc + 1;
208:    else if key = '↓' then
209:       loc ← loc − 1;
210:    end if
211:    breakpoint;
212: end while
```

(b) ... instrumented to run symbolically

Fig. 1. Example game client.

condition. Having a concrete input that reliably reproduces an error is a great help when trying to correct the bug in the source code.

4.2. Generating Constraints

The first step of our technique is identifying the main event loop of the game client and all of its associated client state, which should include any global memory, memory that is a function of the client input, and memory that holds data received from the network. These state variables are then provided to the symbolic execution tool, which is used to generate a constraint for each path through the loop in a single round. These constraints are thus referred to as *round constraints*.

For example, consider the toy game client in Figure 1(a). This client reads a keystroke from the user and either increments or decrements the value of the location variable *loc* based on this key. The new location value is then sent to the server, and the client loops to read a new key from the user. Although this example is a toy, one can imagine it forming the basis for a *Pong* client.

To prepare for symbolic execution, we modify the program slightly, as shown in Figure 1(b). First, we initialize the variable *key* not with a concrete input value read from the user (line 103) but instead as an unconstrained symbolic variable (line 203). We then replace the instruction to send output to the server (line 111) with a breakpoint in the symbolic execution (line 211). Finally, we create a new symbolic state variable, *prev_loc* (line 200), which will represent the game state up to this point in the execution. The state variable *loc* will be initialized to this previous state (line 201).

Symbolically executing one loop iteration of this modified program, we see that there are four possible paths that the client could take in any given round. In the first possible path, *key* is ESC, and the game ends. Note that this branch never reaches the breakpoint. The second and third possible paths are taken when *key* is equal to ↑ and ↓, respectively. The final path is taken when *key* is none of the aforementioned keys. These last three paths all terminate at the breakpoint.

Via symbolic execution, the verifier can obtain the constraints for all symbolic variables at the time each path reached the breakpoint. Because we artificially created *prev_loc* during the instrumentation phase, it remains an unconstrained symbolic variable in all three cases. The state variable *loc*, however, is constrained differently on each of the three paths. In the case when *key* is equal to ↑, symbolic execution reports $loc = prev\_loc + 1$ as the only constraint on *loc*. When *key* is equal to ↓, the constraint is that $loc = prev\_loc - 1$. And when *key* is not ↑, ↓, or ESC, the constraint is that $loc = prev\_loc$.

Therefore, there are three possible paths that can lead to a message being sent to the server. If the server receives a message from a client—and the client is a valid client—then the client must have taken one of these three paths. Since each path introduces a constraint on the value of *loc* as a function of its previous value, the verifier can take the disjunction of these constraints, along with the current and previous values of *loc* (which the server already knows) and see if they are all logically consistent. That is, the verifier can check to see if the change in values for *loc* match up to a possible path that a valid game client might have taken. If so, then this client is behaving according to the rules of a valid game client. The disjunction of round constraints in this case is:

$$(loc = prev\_loc + 1) \vee (loc = prev\_loc - 1) \vee (loc = prev\_loc). \quad (1)$$

For example, suppose the verifier knows that the client reported on its previous turn that its *loc* was 8. If the client were to then report its new location as $loc = 9$, the verifier could simply check to see if the following is satisfiable:

$$(prev\_loc = 8) \wedge (loc = 9) \wedge$$
$$[(loc = prev\_loc + 1) \vee (loc = prev\_loc - 1) \vee (loc = prev\_loc)].$$

Of course, it is satisfiable, meaning that the new value $loc = 9$ could in fact have been generated by a valid game client. Suppose, though, that in the next turn, the client reports his new position at $loc = 12$. Following the same algorithm, the verifier would check the satisfiability of

$$(prev\_loc = 9) \wedge (loc = 12) \wedge$$
$$[(loc = prev\_loc + 1) \vee (loc = prev\_loc - 1) \vee (loc = prev\_loc)].$$

Because these round constraints are *not* satisfiable, no valid game client could have produced the message $loc = 12$ (in this context). Therefore, the verifier can safely conclude that the sender of that message is running an incompatible game client—is cheating.

There are also constraints associated with the variable *key*. We have omitted these here for clarity, showing only the constraints on *loc*. We have also omitted the constraints generated by the preamble of the loop, which in this case are trivial ("$loc = 0$") but in general would be obtained by applying symbolic execution to the preamble separately. Had there been any random coin flips or reading of the current time, the variables storing the results would also have been declared symbolic, and constraints generated accordingly. While file input (e.g., configuration files) could also be declared symbolic, in this paper we generally assume that such input files are known to the verifier (e.g., if necessary, sent to the server at the beginning of game play) and so treat these as concrete.

4.3. Accumulating Constraints

While the branches taken by a client in each round may not be visible to the verifier, the verifier can keep a set of constraints that represent possible client executions so far. Specifically, the verifier forms a conjunction of round constraints that represents a sequence of possible paths through the client's loop taken over multiple rounds; we call this conjunction an *accumulated constraint* and denote the set of satisfiable accumulated constraints at the end of round i by $\mathcal{C}_i$. This set corresponds to the possible paths taken by a client through round i.

The verifier updates a given set $\mathcal{C}_{i-1}$ of accumulated constraints upon receiving a new client message msg_i in round i. To do so, the verifier first combines the values given in msg_i with each round constraint for round i, where each symbolic variable in the round constraint represents client state for round i, and the round constraint characterizes

```
300: C_i ← ∅
301: M ← msgToConstraint(msg_i)
302: for G ∈ G(i) do
303:    for C ∈ C_{i−1} do
304:       C′ ← C ∧ G ∧ M
305:       if isSatisfiable(C′) then
306:          C_i ← C_i ∪ {C′}
307:       end if
308:    end for
309: end for
```

Fig. 2. Construction of $\mathcal{C}_i$ from $\mathcal{C}_{i-1}$ and msg_i.

those variables as a function of the variables for round $i-1$. The verifier then combines each result with each accumulated constraint in $\mathcal{C}_{i-1}$ and checks for satisfiability.

For example, let us parameterize the round constraints for the toy example in Section 4.2 with the round number j:

$$\mathcal{G}(j) = \{loc_j = loc_{j-1} + 1\ ,\ loc_j = loc_{j-1} - 1\ ,\ loc_j = loc_{j-1}\}.$$

Note that each member of $\mathcal{G}(j)$ corresponds to a disjunct in (1). If in round $i = 2$ the server receives the message $msg_2 = 9$ from the client, then it generates the constraint $M =$ "$loc_2 = 9$", because the value "9" in the message represents information corresponding to the variable *loc* in the client code. Then, combining M with each $G \in \mathcal{G}(2)$ gives the three constraints:

$$\begin{aligned} &loc_2 = 9 \wedge loc_2 = loc_1 + 1 \\ &loc_2 = 9 \wedge loc_2 = loc_1 - 1 \\ &loc_2 = 9 \wedge loc_2 = loc_1. \end{aligned}$$

Note that the combination of the client message with each round constraint involves both instantiation (e.g., using $j = 2$ above) as well as including the specific values given in the client message at that round (i.e., $loc_2 = 9$ above).

These three round constraints each represent a possible path the client might have taken in the second round. The verifier must therefore consider each of them in turn as if it were the correct path. For example, if $\mathcal{C}_1 = \{loc_1 = 8\}$, then the verifier can use each round constraint to generate the following possible accumulated constraints:

$$\begin{aligned} &loc_1 = 8 \wedge [loc_2 = 9 \wedge loc_2 = loc_1 + 1] \\ &loc_1 = 8 \wedge [loc_2 = 9 \wedge loc_2 = loc_1 - 1] \\ &loc_1 = 8 \wedge [loc_2 = 9 \wedge loc_2 = loc_1] \end{aligned}$$

Since the second and third constraints are not satisfiable, however, this reduces to

$$\begin{aligned} \mathcal{C}_2 &= \{loc_1 = 8 \wedge [loc_2 = 9 \wedge loc_2 = loc_1 + 1]\} \\ &= \{loc_1 = 8 \wedge loc_2 = 9\} \end{aligned}$$

The basic algorithm for constructing $\mathcal{C}_i$ from $\mathcal{C}_{i-1}$ and msg_i is thus as shown in Figure 2. In this figure, msgToConstraint simply translates a message to the constraint representing what values were sent in the message. It is important to note that while $|\mathcal{C}_i| = 1$ for each i in our toy example, this will not generally be the case for a more complex game. In another game, there might be many accumulated constraints represented in $\mathcal{C}_{i-1}$, each of which would have to be extended with the possible new round constraints to produce $\mathcal{C}_i$.

4.4. Constraint Pruning

Every accumulated constraint in $\mathcal{C}_i$ is a conjunction $C = c_1 \wedge \ldots \wedge c_n$ (or can be written as one, in conjunctive normal form). In practice, constraints can grow very quickly. Even

in the toy example of the previous section, the accumulated constraint in C_2 has one more conjunct than the accumulated constraint in C_1. As such, the verifier must take measures to avoid duplicate constraint checking and to reduce the size of accumulated constraints.

First, the verifier partitions the conjuncts of each new accumulated constraint C' (line 304) based on variables (e.g., loc_2) referenced by its conjuncts. Specifically, consider the undirected graph in which each conjunct c_k in C' is represented as a node and the edge $(c_k, c_{k'})$ exists if and only if there is a variable that appears in both c_k and $c_{k'}$. Then, each connected component of this graph defines a block in the partition of C'. Because no two blocks for C' share variable references, the verifier can check each block for satisfiability independently (line 305), and each block is smaller, making each such check more efficient. And, since some accumulated constraints C' will share conjuncts, caching proofs of satisfiability for previously checked blocks will allow shared blocks to be confirmed as satisfiable more efficiently.

Second, because round constraints refer only to variables in two consecutive rounds; that is, any $G \in \mathcal{G}(j)$ refers only to variables for round j and $j-1$—the formulas G and M in line 304 will refer only to variables in rounds i and $i-1$. Therefore, if there are blocks of conjuncts for C' in line 304 that contain no references to variables for round i, then these conjuncts cannot be rendered unsatisfiable in future rounds. Once the verifier determines that this block of conjuncts is satisfiable (line 305), it can safely remove the conjuncts in that block from C'.

4.5. Server Messages

Round constraints are not a function of only user inputs (and potentially random coin flips and time readings) but also of messages from the server that the client processes in that round. We have explored two implementation strategies for accounting for server messages when generating round constraints.

—*Eager*. In this approach, *eager round constraints* are generated with the server-to-client messages marked symbolic in the client software, just like user inputs. Each member of $\mathcal{G}(i)$ is then built by conjoining an eager round constraint with one or more conjuncts of the form "$svrmsg = m$", where *svrmsg* is the symbolic variable for a server message in the client software, and m is the concrete server message that this variable took on in round i. We refer to this approach as "eager" since it enables precomputation of round constraints prior to verification but, in doing so, also computes them for paths that may never be traversed in actual game play.

—*Lazy*. In this approach, *lazy round constraints* are generated from the client software after it has been instantiated with the concrete server-to-client messages that the client processed in that round; these round constraints for round i then constitute $\mathcal{G}(i)$ directly. Since the server messages are themselves a function of game play, the lazy round constraints cannot be precomputed (as opposed to eager round constraints) but rather must be computed as part of verification. As such, the expense of symbolic execution is incurred during verification, but only those paths consistent with server messages observed during game play need be explored.

In either case, it is necessary that the server log the messages it sent and that the verifier know which of these messages the client actually processed (versus, say, were lost). In our case study in Section 5, we will discuss how we convey this information to the server, which it records for the verifier.

As discussed above, the eager approach permits symbolic execution to be decoupled from verification, in that eager round constraints can be computed in advance of game play and then augmented with additional conjuncts that represent server messages processed by the client in that round. As such, the generation of round constraints

in the eager approach is a conceptually direct application of a tool like KLEE (albeit one fraught with game-specific challenges, such as those we discuss in Section 5.4.1). The lazy approach, however, tightly couples the generation of round constraints and verification; below we briefly elaborate on its implementation.

To support the lazy approach, we extend KLEE by building a model of the network that permits it access to the log of messages the client processed (from the server) in the current round i and any message the client sent in that round. Below, we use the term *active path* to refer to an individual, symbolically executing path through the client code. Each active path has its own index into the message log, so that each can interact with the log independently.

To handle server-to-client messages from the log, we intercept the `recv()` system call and instead call our own replacement function. This function first checks to see that the next message in the network log is indeed a server-to-client message. If it is, we return the message and advance this active path's pointer in the log by one message. Otherwise, this active path has attempted more network reads in round i than actually occurred in the network log prior to reaching the breakpoint corresponding to a client-message send. In this case, we return zero bytes to the `recv()` call, indicating that no message is available to be read. Upon an active path reaching the breakpoint (which corresponds to a client send), if the next message in the log is a server-to-client message, then this active path has attempted fewer network reads than the log indicates, and it is terminated as invalid. Otherwise, the round constraint built so far is added to $\mathcal{G}(i)$, and the logged client message is used to instantiate the new conjunct M in line 301 of Figure 2.

4.6. Scaling to Many Clients

Implementing our technique on a real-world online game with a large user base might require its own special implementation considerations. As we will see in Section 5, our eager and lazy implementations are not yet fast enough to perform validation on the critical path of game play. So, the game operator must log all the messages to and from clients that are needed to validate game play offline. That said, the need for logging will not be news to game operators, and they already do so extensively.

> LOG EVERYTHING, and offer a robust system for reviewing the logs. When hunting down bugs and/or reviewing player cries of foul, nothing makes the job of the GM easier than knowing that he/she has perfect information and can state with 100% accuracy when a player isn't telling the whole truth.
> —D. Schubert [Mulligan and Patrovsky 2003, p. 221]

As such, our approach introduces potentially little additional logging to what game operators already perform. Nevertheless, to minimize this overhead, game operators might use a log-structured file system [Rosenblum and Ousterhout 1992], which is optimized for small writes (as would be the case when logging client and server messages). Log-structured file systems have been implemented for NetBSD and Linux, for example.

Once the messages are logged, they can be searched later to extract a specific game trace to be checked (e.g., for a winning player). The checking itself can be parallelized extensively, in that the trace of a player can be checked independently of others', and even blocks within accumulated constraints C' (see Section 4.4) can be checked in parallel. Traces can also be partially checked, by starting in the middle of a trace, say at round i with client-to-server message msg_i, and checking from that point forward (i.e., with $\mathcal{C}_{i-1} = \{\text{true}\}$). Of course, while such a partial check can validate the internal consistency of the part of the trace that is checked, it will not detect inconsistencies between the validated part and other parts.

5. CASE STUDY: *XPILOT*

In our first case study, we apply our technique to *XPilot*, an open-source multiplayer game written in about 150,000 lines of C code. *XPilot* uses a client-server architecture that has influenced other popular open source games. For example, the authors of *Freeciv* used *XPilot*'s client-server architecture as a basis for the networking in that game. *XPilot* was first released over 15 years ago, but it continues to enjoy an active user base. In fact, in July 2009, 7b5 Labs released an *XPilot* client for the Apple iPhone and Apple iPod Touch,[1] which is one of several forks and ports of the *XPilot* code base over the years. We focus on one in particular called *XPilot NG* (*XPilot Next Generation*).

5.1. The Game

The game's style resembles that of *Asteroids*, in which the player controls an avatar in the form of a spaceship, which she navigates through space, avoiding obstacles and battling other ships. But *XPilot* adds many new dimensions to game play, including computer-controlled players, several multiplayer modes (capture the flag, death match, racing, etc.), networking (needed for multiplayer), better physics simulation (e.g., accounting for fuel weight in acceleration), and updated graphics. In addition, *XPilot* is a highly configurable game, both at the client and the server. For example, clients can set key mappings, and servers can configure nearly every aspect of the game (e.g., ship mass, initial player inventory, probability of each type of power-up appearing on the map, etc.).

As we have discussed, developers of today's networked games design clients with little authoritative state in order to help address cheating. In keeping with that paradigm, *XPilot* was written with very little such state in the client itself. Despite this provision, there are still ways a malicious user can send invalid messages in an attempt to cheat. In *XPilot*, there are some sets of keys that the client should *never* report pressing simultaneously. For example, a player cannot press the key to fire (`KEY_FIRE_SHOT`) while at the same time pressing the key to activate his shield (`KEY_SHIELD`). A valid game client will filter out any attempts to do so, deactivating the shield whenever a player is firing and bringing it back online afterward. However, an invalid game client might attempt to gain an advantage by sending a keyboard update that includes both keys. As it happens, the server does its own (manually configured) checking and so the cheat fails in this case, but the fact that the client behavior is verifiably invalid remains. There are numerous examples of similar cheats in online games that servers fail to catch, either because of programming errors or because those particular misuses of the protocol were unforeseen by the game developers. In our evaluations, we confirmed that our technique detects this attempt to cheat in *XPilot*, as expected. This detection was a direct result of the logic inherent in the game client, in contrast to the manually programmed rule in the *XPilot* server. While this example is illustrative, we emphasize that our goal is not to identify new cheating vulnerabilities on the *XPilot* server, but rather to illustrate how a pre-existing game client can be adapted for verification in our framework and how the verifier performs under different configurations.

At the core of the architecture of the *XPilot* client is a main loop that reads input from the user, sends messages to the server, and processes new messages from the server. In Sections 5.3 and 5.4, we describe the verification of *XPilot* client behavior by generating lazy round constraints and eager round constraints for this loop, respectively. However, we first describe modifications we made to *XPilot*, in order to perform verification.

[1] http://7b5labs.com/xpilotiphone.

5.2. Game Modifications

Message acknowledgments. Client-server communication in *XPilot* uses UDP traffic for its timeliness and decreased overhead; the majority of in-game packets are relevant only within a short time after they are sent (e.g., information about the current game round). For any traffic that must be delivered reliably (e.g., chat messages between players), *XPilot* uses a custom layer built atop UDP. Due to *XPilot*'s use of UDP and the fact that it can process arbitrary numbers of messages in a single client loop, we added to *XPilot* an acknowledgement scheme to inform the server of which inbound messages the client processed in each loop iteration and between sending its own messages to the server. The server logs this information for use by the verifier. There are many possible efficient acknowledgement schemes to convey this information; the one we describe in Appendix A assumes that out-of-order arrival of server messages is rare.

These acknowledgments enable the server to record a log of relevant client events in the order they happened (as reported by the client). For each client-to-server message that the server never received, the verifier simply replaces the constraint M implied by the missing message (see line 301 of Figure 2) with $M = \mathsf{true}$.

Floating-point operations. *XPilot*, like most games of even moderate size, includes an abundance of floating-point variables and math. However, it is not currently possible to generate constraints on floating-point numbers with KLEE or to check them using STP. Therefore, we implement *XPilot*'s floating-point operations using a simple fixed-point library of our own creation. As a result, symbolic execution on the *XPilot* client produces constraints from this library for every mathematical operation in the client code involving a symbolic floating-point number. These constraints, in turn, inflate the verification speeds reported in Section 5.4, in particular.

Bounding loops. The number of round constraints can grow rapidly as new branch points are encountered during path traversal. Loops in the code can be especially problematic; a loop with up to n iterations induces $\Omega(n^2)$ round constraints. During symbolic execution of *XPilot*, most loops have a concrete number of iterations, but there are some loops that iterate over a symbolic variable. If this symbolic variable is unbounded, then the number of round constraints can become impractical to manage. While some loops are not explicitly bounded in the code, they are implicitly bounded by the environment during normal execution. For example, the user input loop in *XPilot* is a while loop that continues until there are no longer any queued user input events to process. During normal execution, the number of iterations of this loop is limited by how fast a player can press a key. As such, during symbolic execution, we limited the input processing loop to a maximum of three iterations because we observed that during gameplay, this queue never contained more than three events. The user input loop and the packet processing loop in *XPilot* required this type of modification, but all other loops were exhaustively searched.

Client trimming. The *XPilot* client, like presumably any game client, contains much code that is focused on enhancing the user gaming experience but that has no effect on the messages that the client could send to the server. To avoid analyzing this code, we trimmed much of it from the game client that we subjected to analysis. Below we summarize the three classes of such code that we trimmed. Aside from these three types of code, we also trimmed mouse input-handling code, since all game activities can be performed equivalently using the keyboard.

First, several types of user inputs impact only the graphical display of the game but have no effect on the game's permissible behaviors as seen by the server. For example, one type of key press adjusts the display of game-play statistics on the user's console. As such, we excised these inputs from the client software for the purposes of our analysis.

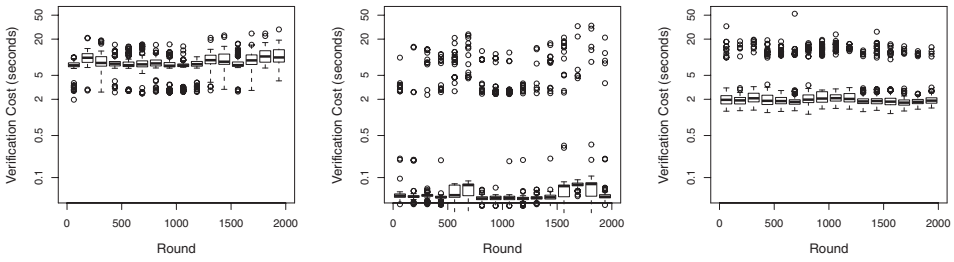

(a) Cost per round (lazy)

(b) Cost per round (lazy) with *XPilot*-specific optimizations

(c) Cost per round (eager) with *XPilot*-specific optimizations

Fig. 3. Verification cost per round while checking a 2,000-round *XPilot* game log.

Second, there are certain "reliable" messages the server sends the client (using the custom reliable-delivery protocol built over UDP). Reliable traffic is vital to the set-up and tear-down of games and game connections, but once play has begun, reliable messages are irrelevant for game play. Types of messages the server sends reliably are in-game chat messages (both among players and from the server itself), information about new players that have joined, and score updates, all of which are relatively infrequent and purely informational, in the sense that their delivery does not alter the permissible client behaviors. As such, we ignored them for the purpose of our analysis.

Third, KLEE is built upon LLVM and requires the input executable to be compiled into the LLVM intermediate representation (IR). Like all software, *XPilot* does not execute in isolation and makes use of external libraries; not all of these were compiled into LLVM IR. Specifically, the graphics library was not symbolically executed by KLEE, and instead any return values from graphics calls that *XPilot* later needed were simply declared symbolic.

5.3. Verification with Lazy Round Constraints

In this section we measure the performance of verification using lazy round constraints. As discussed in Section 4, lazy round constraints are generated once the client-to-server and server-to-client messages are known. Thus, the only unknown inputs to the game client when generating lazy round constraints are the user inputs and time readings (and random coin flips, but these do not affect server-visible behavior in *XPilot*).

In generating lazy round constraints, we departed slightly from the description of our approach in Section 4, in that we inserted multiple breakpoints in the client event loop, rather than only a single breakpoint. Each breakpoint provides an opportunity to prune accumulated constraints and, in particular, to delete multiple copies of the same accumulated constraint. This is accomplished using a variant of the algorithm in Figure 2, using constraints derived from prefixes of the loop leading to the breakpoint, in place of full round constraints. Some of these extra breakpoints correspond to the (multiple) send locations in *XPilot*'s loop. Aside from this modification, we implemented our approach as described in Section 4.

We ran our lazy client verifier on a 2,000-round *XPilot* game log (about a minute of game-play time) on a single core of a 3GHz processor. Figure 3(a) describes the per-round validation cost (in seconds) using a box-and-whiskers plot per 125 rounds: the box illustrates the 25th, 50th, and 75th percentiles; the whiskers cover points within 1.5 times the interquartile range; and circles denote outliers. The per-round verification times averaged 8.6s with a standard deviation of 2.74s. In these experiments, the verifier's memory usage remained below 256MB. As an aside, in every

round, there was exactly one remaining satisfiable accumulated constraint, indicating that, without client state, there is little ambiguity at the verifier about exactly what is happening inside the client program, even from across the network.

By employing an *XPilot*-specific optimization, we were able to significantly improve verification performance. After the trimming described in Section 5.2, the user input paths that we included within our symbolic execution of the client each caused another client-to-server message to be sent, and so the number of such sends in a round indicates to the verifier an upper bound on the number of user inputs in that round. As such, we could tune the verifier's symbolic execution to explore only paths through the client where the number of invocations of the input-handling function equals the number of client messages for this round in the log. This optimization yields the graph in Figure 3(b). Notice that there are three distinct bands in the graph, corresponding to how many times the input-handling function within the game client was called. The first band contains rounds which called the input handler zero times and represents the majority (90.1%) of the total rounds. These rounds were the quickest to process, with a mean cost of 53.8ms and a standard deviation of 21.1ms. The next-largest band (5.1%) contains rounds which called the input handler only once. These rounds took longer to process, with a mean of 3.26s and a standard deviation of 1.05s. The final band represents rounds with more than one call to the input-handling function. This band took the longest to process (12.9s, on average), but it was also the smallest, representing only 4.1% of all rounds.

5.4. Verification with Eager Round Constraints

In this section we discuss verification of *XPilot* using eager constraint generation. Recall that eager round constraints are precomputed from the sanctioned client software without knowledge of the messages the client will process in any given loop iteration. However, we found this approach to require moderate manual tuning to be practical, as we describe below.

5.4.1. Manual Tuning. A direct application of our method for generating eager round constraints for the *XPilot* client loop would replace the user key press with symbolic input and any incoming server message with a symbolic buffer and then use KLEE to symbolically execute the resulting client program. Such a direct application, however, encountered several difficulties. In this section we describe the main difficulties we encountered in this direct approach and the primary adaptations that we made in order to apply it to the *XPilot* client. These adaptations highlight an important lesson: the eager technique, while largely automatic, can require some manual tuning to be practical. Because our technique is targeted toward game developers, we believe that allowing for such manual tuning is appropriate.

Frame processing. In *XPilot*, messages from the server to the client describing the current game state are called *frames*. Each frame is formed of a chain of game *packets* (not to be confused with network packets). The first and last packets in a frame are always special start-of-frame and end-of-frame packets, called PKT_START and PKT_END. Figure 4 shows an *XPilot* frame, containing a packet of type PKT_FUEL and potentially others (indicated by "..."). Packets are encoded as a single header byte followed by a packet data section that can carry anything from a single byte to an arbitrary-length string, depending on the packet type. Frames may contain multiple packet types and multiple instances of the same packet type.

Consider the client's frame-processing algorithm. Given a frame, it first reads the packet header (i.e., the first byte), then calls the handler for that packet, which processes the packet and advances the frame pointer so that the new "first byte" is the packet header of the next packet in the frame. This continues until the packet handler

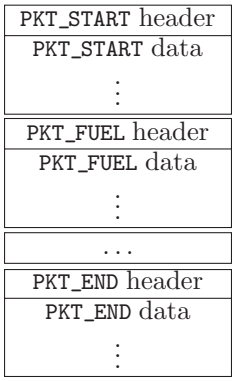

Fig. 4. *XPilot* frame layout.

for PKT_END is called, the return of which signifies the end of the frame handling. Therefore, given a completely symbolic buffer representing the frame, our symbolic execution would need to walk the client code for each possible sequence of packets in a frame, up to the maximum frame size. But *XPilot* has dozens of packet types, some of which include a very small amount data. As evidence of the infeasibility of such an approach, consider the following (very conservative) lower bound on the number of packet sequences: There are at least 10 types of packets that we considered whose total size is at most 5 bytes. The maximum size for a server-to-client frame in *XPilot* is 4,096 bytes, which means there is room for over 800 of these packets. That gives *at least* 10^{800} possible packet sequences that symbolic execution would traverse to generate constraints, which is obviously infeasible.

To make eager constraint generation feasible, then, we adapt our approach to generate round constraints by starting and stopping symbolic execution at multiple points within the loop, as opposed to just the beginning and end. In particular, we apply symbolic execution to the frame-processing and user input-processing portions of the loop separately, to obtain *user-input constraints* and *frame-processing constraints*, which in turn the verifier pieces together during verification to construct the round constraints. Moreover, the verifier can construct the frame-processing constraints on the basis of the particular frame the server sent to the client. It does so dynamically from packet-processing constraints that characterize how the client should process each packet in the particular frame. For example, if the only packet types were PKT_START, PKT_FUEL, PKT_TIME_LEFT, and PKT_END, the packet-processing constraints representing the processing of a single packet would be

$$(p = \texttt{PKT\_START}) \wedge (\text{constraints\_for}(\texttt{PKT\_START}))$$
$$(p = \texttt{PKT\_FUEL}) \wedge (\text{constraints\_for}(\texttt{PKT\_FUEL}))$$
$$(p = \texttt{PKT\_TIME\_LEFT}) \wedge (\text{constraints\_for}(\texttt{PKT\_TIME\_LEFT}))$$
$$(p = \texttt{PKT\_END}) \wedge (\text{constraints\_for}(\texttt{PKT\_END})),$$

where p is a variable for the packet type and constraints_for(PKT_START) represents the additional constraints that would result from symbolic execution of the packet handler for PKT_START. With this new model of packet processing, the verifier can build a frame-processing constraint to represent any given frame from the logs. In this way, when the verifier checks the behavior of a given client, it does so armed with the frames the server sent to the client, the messages the server received from the client, and the frame-processing constraints that characterize the client's processing of each frame, which the verifier constructs from the packet-processing constraints.

Packet processing. Certain individual packet types present their own tractability challenges as well. For example, the payload for a certain packet begins with a 32-bit mask followed by one byte for each bit in the mask that is equal to 1. The client then stores these remaining bytes in a 32-byte array at the offsets determined by the mask (setting any bytes not included in the message to 0). In the packet handler, the *XPilot* client code must sample the value of each bit in the mask in turn. Since the payload (and thus the mask) is symbolic, each of these conditionals results in a fork of two separate paths (for the two possible values of the bit in question). Our symbolic execution of this packet handler, then, would produce over 4 billion round constraints, which is again infeasible. We could have changed the *XPilot* protocol to avoid using the mask, sending 32 bytes each time, but doing so would increase network bandwidth needlessly. Instead, we note that the result of this packet handler is that the destination array is set according to the mask and the rules of the protocol. We thus added a simple rule to the verifier that, when processing this type of packet, generates a constraint defining the value of the destination array directly, as the packet handler would have. Then, when symbolically executing the packet handlers, we can simply skip this packet.

To avoid similar modifications to the extent possible, we pruned the packets the verifier considers during verification to only those that are necessary. That is, there are several packet types that will not alter the permissible behaviors of the client as could be witnessed by the server, and so we ignored them when applying our technique. Most of these packet types represent purely graphical information. For example, a packet of type `PKT_ITEM` simply reports to the client that a game item of a given type (e.g., a power-up or a new weapon) is floating nearby at the given coordinates. This information allows the client to draw the item on the screen, but it does not affect the valid client behaviors as observable by the verifier.[2]

User input. The first part of the client input loop checks for and handles input from the player. Gathering user-input constraints is fairly straightforward, with the exception that *XPilot* allows players to do an extensive amount of keyboard mapping, including configurations in which multiple keys are bound to the same function, for example. We simplified the generation of constraints by focusing on the user actions themselves rather than the physical key presses that caused them. That is, while generating constraints within the user-input portion of *XPilot*, we begin symbolic execution *after* the client code looks up the in-game action bound to the specific physical key pressed, but *before* the client code processes that action. For example, if a user has bound the action `KEY_FIRE_SHOT` to the key `a`, our analysis would focus on the effects of the action `KEY_FIRE_SHOT`, ignoring the actual key to which it is bound. However, as with other client configuration options, the keyboard mapping could easily be sent to the server as a requirement of joining the game, invoking a small, one-time bandwidth cost that would allow the verifier to check the physical key configuration.

5.4.2. Eager Verification Performance. We ran our eager client verifier on the same 2,000-round *XPilot* game log and on the same computer used in Section 5.3. Figure 3(c) describes the per-round validation cost (in seconds) using a box-and-whiskers plot. As in Figure 3(b), we employed here an *XPilot*-specific optimization by observing that the number of client messages in a round bounds the number of user inputs in that round. As such, in piecing together round constraints, the verifier includes a number of copies of user-input constraints (see Section 5.4.1) equal to the client sends in that round.

[2]In particular, whether the client processes this packet is irrelevant to determining whether the client can pick up the game item described in the packet. Whether the client obtains the item is unilaterally determined *by the server* based on it computing the client's location using the low-level client events it receives—an example of how nearly all control is stripped from clients in today's games, because they cannot be trusted.

Similar to Figure 3(b), Figure 3(c) exhibits three bands (the third comprising a few large values), corresponding to different numbers of copies. The large percentage of rounds contained no user inputs and were the quickest to process, with a mean cost of 1.64s and a standard deviation of 0.232s. The second band of rounds—those with a single user input—took longer to process, with a mean of 11.3s and a standard deviation of 1.68s. Remaining rounds contained multiple user inputs and took the longest to process (34.2s, on average), but recall that they were by far the least frequent. The verifier's memory usage remained below 100MB throughout these verification runs.

Comparing Figures 3(b) and 3(c), the times for the eager approach are much slower than those for the lazy approach, when applied to *XPilot*. This performance loss is due to the fact that a large portion of the *XPilot* client code is dedicated to handling server messages. And while the verifier in the eager case has preprocessed this portion of the code, the resulting round constraints are much more complex than in the lazy approach, where the verifier knows the exact values of the server messages when generating round constraints. This complexity results in constraint solving in the eager case (line 305 of Figure 2) being more expensive.

It is also important to recall that lazy and eager are not interchangeable, at least in terms of game developer effort. As discussed in Section 5.4.1, achieving feasible generation of eager round constraints required substantial additional manual tuning, and consequently greater opportunity for programmer error. As such, it appears that the eager approach is inferior to the lazy approach for *XPilot*. Another comparison between the two approaches, with differing results, will be given in Section 6.

6. CASE STUDY: *CAP-MAN*

Our client verification technique challenges the current game-design philosophy by allowing servers to relinquish authoritative state to clients while retaining the ability to validate client behavior and thus detect cheating. As a way of demonstrating this notion, we have written a game called *Cap-Man* that is based on the game *Pac-Man*. In some ways *Cap-Man* is easier to validate than *XPilot* was; it represents a considerably smaller code base (roughly 1,000 lines of C code) and state size.

That said, *Cap-Man* is interesting as a case study for three reasons. First, whereas *XPilot* was written with virtually no authoritative client state, we will see that *Cap-Man* is intentionally rife with it, providing a more interesting challenge for our technique because it is so much more vulnerable to invalid messages. Second, the size of its code base allows us to conduct a more direct comparison between lazy and eager verification. Third, *Cap-Man* differs from *XPilot* in that the set of possible user inputs per round is substantially larger than the set of paths through the client's event loop. That is, in *XPilot*, there is nearly a one-to-one correspondence between user inputs and paths through the client event loop, which dampens the improvement that our technique offers over, for instance, the verifier simply running the client on all possible inputs in each round. *Cap-Man* demonstrates the scalability of our technique to many possible user inputs when this is not the case, thereby separating our technique from other such possible approaches.

6.1. The Game

Cap-Man is a *Pac-Man*-like game in which a player controls an avatar that is allowed to move through a discrete, two-dimensional map with the aim of consuming all remaining "food" items before being caught by the various enemies (who are also navigating the map). Each map location is either an impenetrable wall or an open space, and the open spaces can contain an avatar, an enemy, pieces of food, a power-up, or nothing at all. When a player reaches a map location that contains food or a power-up, he automatically consumes it. Upon consuming a power-up, the player enters a temporary

"power-up mode," during which his pursuers reverse course—trying to escape rather than pursue him—and he is able to consume (and temporarily displace) them if he can catch them. In addition to these features (which were present in *Pac-Man* as well), we have added a new feature to *Cap-Man* to invite further abuse and create more uncertainty at the server: A player may set a bomb (at his current location), which will then detonate a number rounds in the future selected by the user from a predefined range (in our implementation, between 3 and 15 rounds).[3] When it detonates, it kills any enemies (or the player himself) within a certain radius on the map. Players are not allowed to set a new bomb until their previous bomb has detonated.

Cap-Man uses a client-server architecture, which we designed specifically to go against current game-development best practices: that is, it is the *server*, not the client, which has a minimum of authoritative state. The client tracks his own map position, power-up-mode time remaining, and bomb-placement details. Specifically, at every round, the client sends a message to the server indicating its current map position and remaining time in power-up mode. It also sends the position of a bomb explosion, if there was one during that round. Note that the client never informs the server when it decides to *set* a bomb. It merely announces when and where detonation has occurred. The server, in contrast, sends the client the updated positions of his enemies, this being the only game state for which the server has the authoritative copy.

The design of *Cap-Man* leaves it intentionally vulnerable to a host of invalid-message attacks. For example, although valid game clients allow only contiguous paths through the map, a cheating player can arbitrarily adjust his coordinates, ignoring the rules of the game: a cheat known in game-security parlance as "telehacking." He might also put himself into power-up mode at will, without bothering to actually consume a power-up. Finally, there is no check at the server to see whether or not a player is lying about a bomb placement by, for example, announcing an explosion at coordinates that he had not actually occupied within the past 15 rounds. In fact, the *Cap-Man* server contains no information about (or manual checks regarding) the internal logic of the game client.

In order to detect cheating in *Cap-Man*, we apply our technique in both its lazy and eager variations. Due to *Cap-Man*'s smaller size and simpler code structure, we can generate round constraints over an entire iteration of the main loop in each case, without the need to compartmentalize the code and adopt significant trimming measures as we did for *XPilot*.

6.2. Evaluation

Using our technique, we are able to detect invalid-command cheats of all the types listed above. Below we present the results of client-validity checks on a game log consisting of 2,000 rounds (about 6–7 minutes of game-play time), during which the player moved around the map randomly, performing (legal) bomb placements at random intervals.

Figure 5 shows that the verification costs for *Cap-Man* were consistently small, with a mean and standard deviation of 752ms and 645ms for verification via lazy round constraints (Figure 5(a)) and a mean and standard deviation of 310ms and 193ms for verification using eager round constraints (Figure 5(b)). The lazy method was (on average) roughly 2.5 times slower than the eager method, owing to the overhead of symbolic execution to compute round constraints for each round individually during verification. The verifier's memory usage remained below 100MB throughout verification of both types. While in the *XPilot* case study, eager verification required significantly greater

[3]In our preliminary work [Bethea et al. 2010], a bomb detonated in a fixed number of rounds. We changed the game to accommodate a user-selected number of rounds to bomb detonation in order to demonstrate the ability of our technique to scale to a larger number of possible user inputs.

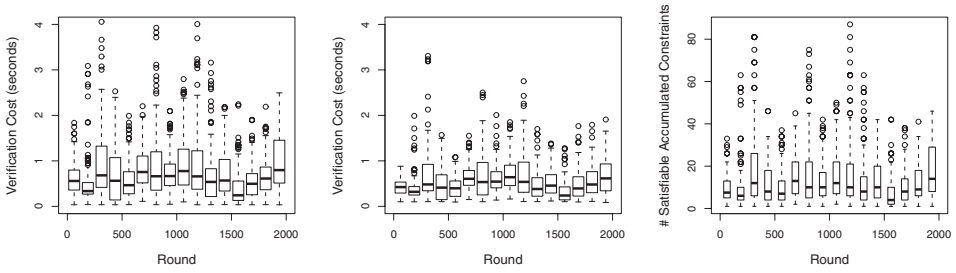

(a) Cost per round (lazy) (b) Cost per round (eager) (c) Satisfiable accumulated constraints per round (eager)

Fig. 5. Verifying a 2,000-round *Cap-Man* game log.

development effort (see Section 5.4.1), this additional effort was unnecessary with *Cap-Man* due to its relative simplicity.

Figure 5(c) shows the number of satisfiable accumulated constraints during eager verification, which did not trend upward during the run. In lazy verification, the number of satisfiable accumulated constraints was virtually identical. (Variations in our pruning implementations caused less than 1% of the rounds to differ, and then by at most 12 accumulated constraints.) In the case of *XPilot*, the number of satisfiable accumulated constraints was always 1, but in *Cap-Man* there were often multiple accumulated constraints that remained satisfiable at any given round. This increase resulted primarily from state the *Cap-Man* client maintains but does not immediately report to the server (e.g., whether a bomb has been set, and with what detonation timer). The relationship between this hidden state and the number of satisfiable accumulated constraints is an important one. Consider the verification of a *Cap-Man* game that is currently in round i, with no bomb placements in the last 15 rounds (unbeknownst to the verifier). The verifier must maintain accumulated constraints that reflect possible bomb placements at each of rounds $i-14$ through i. Upon encountering msg_{i+1} with an announcement of a bomb explosion, the verifier can discard not only all current accumulated constraints which do *not* include a bomb placement in any of rounds $i-14$ through $i-2$, but also those accumulated constraints which *do* include bomb placements in rounds $i-1$ through $i+1$, because players can only have one pending bomb at a time. This rule was not manually configured into the verifier; it was inferred automatically from the client code.

7. CASE STUDY: *TETRINET*

As discussed in Section 1, our verification technique presents opportunities to reduce the bandwidth consumed by an online game, since it allows the client's management of state to be verified with less-than-complete information. In our third case study, we used a preexisting game called *TetriNET* in order to explore simple bandwidth-savings measures and their impact on client verification performance.

TetriNET is a multiplayer clone of the classic puzzle game *Tetris*. It was originally developed in 1997 but remains popular and has been reimplemented on many different platforms. In the game, random *tetrominoes*, which are geometric shapes consisting of four connected blocks, automatically advance down each player's playing field, one at a time. As each does, the player can rotate it into four possible orientations or slide it horizontally left or right. The objective of the game is to arrange tetrominoes so that, as they land, they create gapless horizontal rows of blocks on the playing field. When such a gapless row is created, it is cleared, each row of blocks above it falls down one row, and—in the primary multiplayer departure from *Tetris*—a line with gaps is added

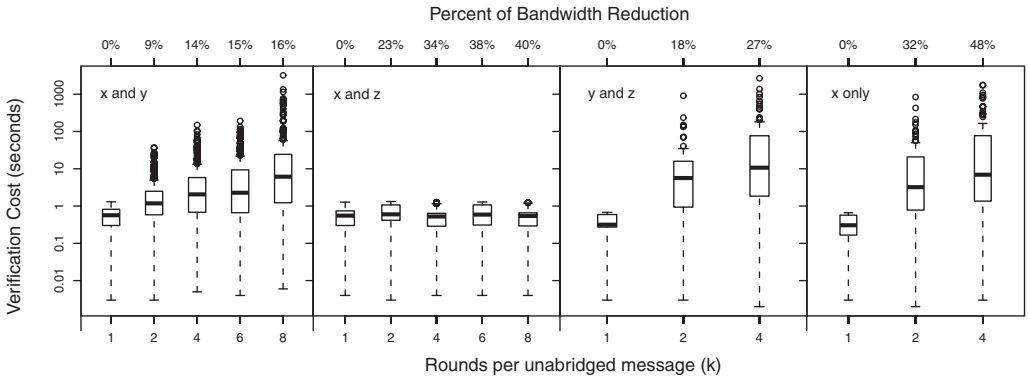

Fig. 6. Verification of 100-round *TetriNET* logs under different configurations of client-to-server message content.

to the bottom of the other players' fields. A player loses when there is no room for an additional tetromino to enter her playing field. *TetriNET* is implemented in C and has a client-server architecture similar to *Cap-Man* and *XPilot*. The server does virtually no checking on client messages and so is vulnerable to cheats; for instance, a client could indicate it cleared a row that has gaps or placed a new tetromino in a spot that the client should not have allowed it to reach.

Enabling the use of our verification tool with *TetriNET* required some consideration of how the user input operates. A tetromino in play automatically moves down one row every 0.5 seconds, and when the piece can move no further, the position is fixed and a new random piece starts falling from the top of the game screen. The placement of the tetromino in a permanent resting point defines the end of single round of gameplay, at which time the x and y coordinates and the rotation z are sent to the server. Until the end of the round, though, the player can move the piece horizontally or rotate it as many times as she wishes. So, even though there is a finite number of final fixed positions for a game piece in a given round, there are theoretically an infinite number of possible inputs sequences that could lead to each valid final position. To minimize the number of input sequences that must be explored symbolically, we considered a restricted version of gameplay where the gameboard must be empty above each tetromino at the time of its placement, and so three rotations and six horizontal moves sufficed to reach any such placement on the 12-column gameboard.

To demonstrate bandwidth reduction in *TetriNET* using our verification technique, we simply reduced the information in the client-to-server messages. *TetriNET* was modified so that only every k rounds was the complete tuple (x, y, z) sent to the server (an "unabridged message"). In other rounds, the client sent a partial tuple, omitting x, y, or z or both y and z. Figure 6 shows the trade-off between the bandwidth reductions accomplished and the costs of verifying client behavior using the lazy client verifier, where the bandwidth reductions were calculated assuming x, y and z are sent in four, five, and two bits, respectively. (The actual *TetriNET* implementation is not engineered for bandwidth reduction and so uses payload space more wastefully.) These graphs each represent five random play sessions of 100 rounds each and the same five play sessions were used for each of the four experiments. Note that $k = 1$ is equivalent to verification of an unmodified game client. In the experiment where y is omitted, the verification cost does not increase because there is no ambiguity as to y's value when x and z are provided. During these verification runs, the verifier's memory usage remained below 512MB.

8. MESSAGE LOSS

Games today must be built to tolerate a range of networking conditions, including occasional message loss. While there are standard approaches to recovering lost messages, such as message retransmission at the transport level (i.e., using TCP) or at the application level, retransmission is avoided in some games for two reasons. First, the importance of some messages diminishes quickly, and so by the time the message would be retransmitted, the utility of doing so is lost. Second, retransmission can introduce overheads that high-performance games cannot tolerate.

Lost server-to-client messages pose little difficulty to our client verification technique; all the verifier requires is to know what server-to-client messages the client processed and when, which can be communicated from the client efficiently (e.g., see Appendix). Lost client-to-server messages pose more difficulty, however. Intuitively, our technique can handle client message loss by instantiating the constraint M for a missing round-i message msg_i to simply M = "true" in Figure 2. However, this has two negative consequences.

First, from the server's (and verifier's) perspective, it is impossible to distinguish a lost message from one the client only pretended to send. This can be used by a cheating client to gain latitude in terms of the behaviors that the verifier will consider legitimate. For example, whenever a power-up appears on the game map, an altered game client could collect it by reporting its player's position at the power-up's location. So as to not be caught by the verifier, the client could alter its state to reflect having sent messages that would have been induced by the player actually moving to that location, even though these messages were never sent and so, from the server's perspective, were lost. Because it is possible for a valid client on a poor network connection to generate indistinguishable behavior, this cheat is not in the class that our verifier detects. Nevertheless, as discussed in Section 2, our techniques are compatible with existing methods that address this type of cheat.

Another consequence of message loss is that the performance of verification can be severely impacted by it. The performance results in Sections 5–6 did not reflect the loss of any client messages; instead, the game logs that we validated included all messages that the client sent. However, in practice message loss causes the accumulated constraints $\mathcal{C}_i$ to grow dramatically, since *any* path through the client that causes a message to be sent is deemed possible in round i. As a result, in experimenting with message loss in *XPilot*, we found that in the face of lost messages, the performance of our technique decays very substantially.

As such, we propose a lightweight scheme to enable our technique to retain its performance in the face of (limited) message loss. Rather than retransmitting messages, our technique communicates a small amount of additional information per client-to-server message to enable the verifier to prune accumulated constraints effectively in the face of message loss. Intuitively, the client remembers the path through its event loop that it traverses in round i and then conveys evidence of this to the server over the next several messages. The server records this evidence for the verifier, which uses it to prune round constraints considered for round i.

There are several ways to instantiate this intuition within our framework. Here we describe one implementation that works well in *XPilot*. In this implementation, the "evidence" that the client conveys to the server for the path it traversed in round i is a hash of the fields of the message it sent in round i that are a function of (only) the path traversed. Rather than send the entire hash in a subsequent message, however, the client "trickles" this hash value to the server, for instance, one bit per message, so that subsequent message losses still enable the server to collect a number of hash bits for each round. After the client's messages are recorded at the server, the verifier

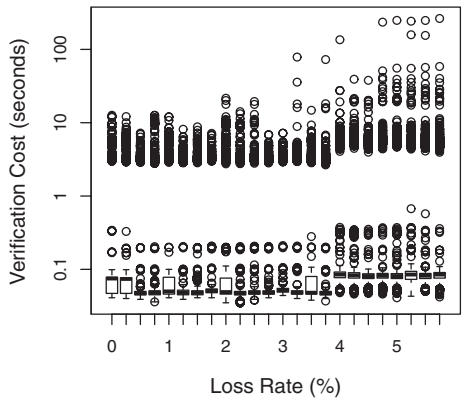

Fig. 7. Verification (lazy) of 2000-round *XPilot* log with loss of client-to-server messages at the rate indicated on horizontal axis.

collects these bits and uses them to prune the round constraints considered at each step of verification where it is missing a message.

We have prototyped this approach in the context of lazy verification, in order to validate the ability of the *XPilot* verifier to retain its performance in the face of message losses. (*Cap-Man* and *TetriNET* use TCP and so do not face message-loss issues.) The hash we use is a 16-bit BSD `sum`, and the k-th bit of the round-i message hash is carried on the round $i + k$ message ($1 \leq k \leq 16$). As such, each message carries an extra 16 bits composed of bits from the previous 16 client-to-server messages.

To show the effectiveness of this approach, we repeated the lazy verification of 2000-round *XPilot* game logs using *XPilot*-specific optimizations (cf. Figure 3(b)) but introduced client-to-server message losses to show that our approach tolerates them seamlessly. We experimented with two types of message loss. In the first, each client-to-server message is lost with a fixed probability. Figure 7 shows, box-and-whiskers plots that illustrate the per-round verification costs that resulted, as a function of this loss rate. Note that a message loss rate of 4% earns a "critical" designation at a real-time monitoring site like www.internetpulse.net. As Figure 7 shows, our technique can easily handle such a high loss rate.

A second type of loss with which we experimented is a burst loss, that is, the loss of a contiguous sequence of client-to-server messages. Figure 8 shows the verification costs per round in five different message logs in which a burst loss of length 6, 10, or 14 client-to-server messages is introduced at a random point between the 100th and 150th round. As these graphs show, the verification costs do tend to spike in the region where the burst loss occurs, but the verification costs remain feasible and recover after the burst to their original durations. Only when the burst length exceeds 16 (not shown) do the verification costs become and remain too large to be practical.

9. CONCLUSION

The need to detect cheats has heavily influenced the design of online games. Cheating has driven game developers to minimize or eliminate authoritative state from game clients. These measures have direct impact on the game operator's bottom line, due to the inflated bandwidth costs that result and to the manual and heuristic (and hence ongoing) effort of programming server-side checks on client behaviors.

In this article we described an approach to validate the server-visible behavior of game clients. Our approach validates that game-client behavior is a subset of the

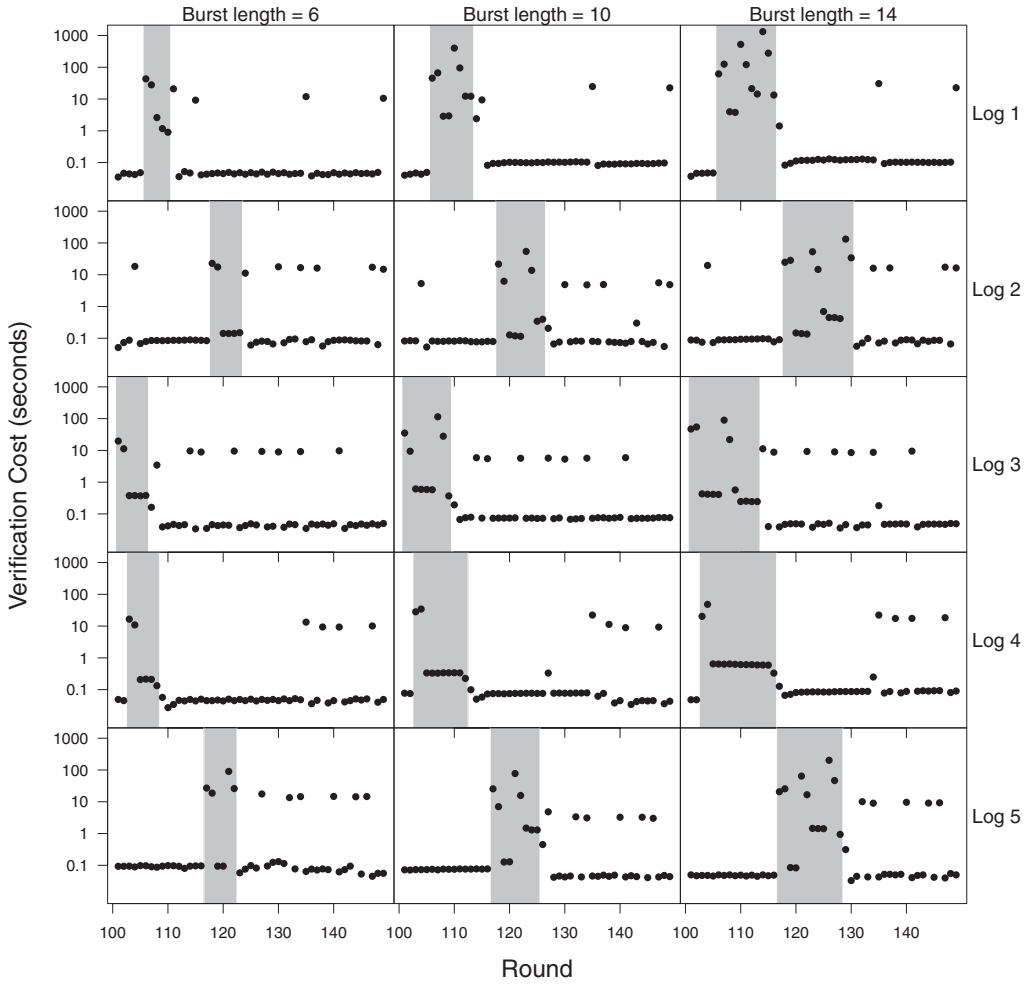

Fig. 8. Verification (lazy) of *XPilot* logs with randomly induced bursts of client-to-server message losses. Shaded areas designate rounds in which losses occurred.

behaviors that would be witnessed from the sanctioned client software, in light of the previous behaviors of the client and the game state sent to that client. Our technique exploits a common structure in game clients, namely a loop that accepts server and user inputs, manages client state, and updates the server with information necessary for multiplayer game play. Our technique applies symbolic execution to this loop to produce constraints that describe its effects. The game operator can then automatically check the consistency of client updates with these constraints offline. We explored both lazy and eager approaches to constraint generation and investigated the programmer effort each entails, as well as their performance.

We demonstrated our technique using three case studies. In the first, we applied our validation approach to *XPilot*, an existing open-source game. We detailed the ways we adapted our technique, in both the lazy and eager variants, to allow for efficient constraint generation and server-side checking. While this effort demonstrated applying our approach to a real game, it was less satisfying as a test for our technique, in that *XPilot* was developed in the mold of modern games, with virtually no authoritative

state at the client. We thus also applied our technique to a simple game of our own design that illustrated the strengths of our technique more clearly. We then showed simple ways to leverage our technique to reduce bandwidth consumption in a game called *TetriNET*, and returned to *XPilot* to demonstrate a strategy for dealing with message loss.

We believe that our technique can change how game developers address an important class of game cheats today, and in doing so opens up new avenues of game design that permit lower bandwidth utilization and better performance.

APPENDIX

A. AN *XPILOT* ACKNOWLEDGEMENT SCHEME

As discussed in Section 5.2, an efficient acknowledgement scheme allows the server (and hence verifier) knowledge of the order (and loop iterations) in which the client processed server messages and sent its own messages. In the following, we describe one such scheme that is optimized for messages that arrive at the client mostly in order.

In this scheme, the *XPilot* client includes a sequence number *c2sNbr* on each message it sends to the server, and similarly the server includes a sequence number *s2cNbr* on each message it sends to the client. Each message from the server to a client also includes the largest value of *c2sNbr* received from that client. In each client message, the client includes *c2sAckd*, the largest value of *c2sNbr* received in a server message so far; a sequence *lateMsgs*[] of server message sequence numbers; and a sequence *eventSeq*[] of symbols that encode events in the order they happened at the client. The symbols in *eventSeq*[] can be any of the following. Below, *s2cAckd* is the largest sequence number *s2cNbr* received by the client before sending message *c2sAckd*, and similarly *loopAckd* is the largest client loop iteration completed at the client prior to it sending *c2sAckd*.

—`Loop` denotes a completed loop iteration. The j-th occurrence of `Loop` in *eventSeq*[] denotes the completion of loop iteration $loopAckd + j$.
—`Send` denotes the sending of a message to the server. The j-th occurrence of `Send` in *eventSeq*[] denotes the sending of client message $c2sAckd + j$.
—`Recv` and `Skip` denote receiving or skipping the next server message in sequence. The j-th occurrence of `Recv` or `Skip` in *eventSeq*[] denotes receiving or skipping, respectively, server message $s2cAckd + j$. Here, a message a skipped if it has not arrived by the time a server message with a larger sequence number arrives, and so a series of one or more `Skip` symbols is followed only by `Recv` in *eventSeq*[].
—`Late` denotes the late arrival of a message, that is, the arrival of a message that was previously skipped. The j-th occurrence of `Late` in *eventSeq*[] denotes the arrival of server message $lateMsgs[j]$.

As such, *lateMsgs*[] contains a sequence number for each server message that arrives after another with a larger sequence number, and so *lateMsgs*[] should be small. *eventSeq*[] may contain more elements, but the symbols can be encoded efficiently, for instance, using Huffman coding [Huffman 1952], and in at most three bits per symbol in the worst case. Note that the server can determine *s2cAckd* and *loopAckd* based on the previous messages received from the client.

ACKNOWLEDGMENTS

We are deeply grateful to Cristian Cadar, Daniel Dunbar, and Dawson Engler for helpful discussions and for permitting us access to an early release of KLEE. Srinivas Krishnan, Alana Libonati, Andy White, and the anonymous reviewers provided helpful comments on drafts of this article.

REFERENCES

ALEXANDER, L. 2008. World of warcraft hits 10 million subscribers. http://www.gamasutra.com/php-bin/news_index.php?story=17062.

BAUGHMAN, N. E. AND LEVINE, B. N. 2001. Cheat-proof playout for centralized and distributed online games. In *Proceedings of IEEE INFOCOM*.

BETHEA, D., COCHRAN, R. A., AND REITER, M. K. 2010. Server-side verification of client behavior in online games. In *Proceedings of the 17th ISOC Network and Distributed System Security Symposium*. 21–36.

BRUMLEY, D., NEWSOME, J., SONG, D., WANG, H., AND JHA, S. 2006. Towards automatic generation of vulnerability-based signatures. In *Proceedings of the IEEE Symposium on Security and Privacy*.

BRUMLEY, D., WANG, H., JHA, S., AND SONG, D. 2007. Creating vulnerability signatures using weakest pre-conditions. In *Proceedings of the Computer Security Foundations Symposium*.

CADAR, C., DUNBAR, D., AND ENGLER, D. 2008. KLEE: Unassisted and automatic generation of high-coverage tests for complex systems programs. In *Proceedings of the 8th USENIX Symposium on Operating Systems Design and Implementation*.

CADAR, C., GANESH, V., PAWLOWSKI, P. M., DILL, D. L., AND ENGLER, D. R. 2006. EXE: Automatically generating inputs of death. In *Proceedings of the 13th ACM Conference on Computer and Communications Security*.

CHEN, K.-T., JIANG, J.-W., HUANG, P., CHU, H.-H., LEI, C.-L., AND CHEN, W.-C. 2006. Identifying MMORPG bots: A traffic analysis approach. In *Proceedings of the ACM SIGCHI International Conference on Advances in Computer Entertainment Technology*.

CHEN, K.-T., PAO, H.-K. K., AND CHANG, H.-C. 2008. Game bot identification based on manifold learning. In *Proceedings of the 7th ACM SIGCOMM Workshop on Network and System Support for Games*. 21–26.

CHONG, S., LIU, J., MYERS, A. C., QI, X., VIKRAM, N., ZHENG, L., AND ZHENG, X. 2007. Secure web applications via automatic partitioning. In *Proceedings of the 21st ACM Symposium on Operating Systems Principles*. 31–44.

CRONIN, E., FILSTRUP, B., AND JAMIN, S. 2003. Cheat-proofing dead reckoned multiplayer games. In *Proceedings of the 2nd International Conference on Application and Development of Computer Games*.

DELAP, M., KNUTSSON, B., LU, H., SOKOLSKY, O., SAMMAPUN, U., LEE, I., AND TSAROUCHIS, C. 2004. Is runtime verification applicable to cheat detection? In *Proceedings of the 3rd ACM SIGCOMM Workshop on Network and System Support for Games*.

FENG, W., KAISER, E., AND SCHLUESSLER, T. 2008. Stealth measurements for cheat detection in on-line games. In *Proceedings of the 7th ACM SIGCOMM Workshop on Network and System Support for Games*. 15–20.

GAMASUTRA STAFF. 2009. Analyst: Online games now $11b of $44b worldwide game market. http://www.gamasutra.com/php-bin/news_index.php?story=23954.

GANESH, V. AND DILL, D. L. 2007. A decision procedure for bit-vectors and arrays. In *Proceedings of the 19th International Conference on Computer Aided Verification (CAV'07)*. 519–531.

GIFFIN, J. T., JHA, S., AND MILLER, B. P. 2002. Detecting manipulated remote call streams. In *Proceedings of the 11th USENIX Security Symposium*.

GOODMAN, J. AND VERBRUGGE, C. 2008. A peer auditing scheme for cheat elimination in MMOGs. In *Proceedings of the 7th ACM SIGCOMM Workshop on Network and System Support for Games*. 9–14.

GUHA, A., KRISHNAMURTHI, S., AND JIM, T. 2009. Using static analysis for Ajax intrusion detection. In *Proceedings of the 18th International World Wide Web Conference*. 561–570.

HOGLUND, G. AND MCGRAW, G. 2007. *Exploiting Online Games: Cheating Massively Distributed Systems*. Addison-Wesley Professional.

HUFFMAN, D. A. 1952. A method for the construction of minimum-redundancy codes. *Proc. Institute Radio Engin. 40,* 9, 1098–1101.

IZAIKU, T., YAMAMOTO, S., MURATA, Y., SHIBATA, N., YASUMOTO, K., AND ITO, M. 2006. Cheat detection for MMORPG on P2P environments. In *Proceedings of the 5th ACM SIGCOMM Workshop on Network and System Support for Games*.

JAGER, I. AND BRUMLEY, D. 2010. Efficient directionless weakest preconditions. Tech. rep. CMU-CyLab-10-002, Cylab, Carnegie Mellon University.

JHA, S., KATZENBEISSER, S., SCHALLHART, C., VEITH, H., AND CHENNEY, S. 2007. Enforcing semantic integrity on untrusted clients in networked virtual environments (extended abstract). In *Proceedings of the IEEE Symposium on Security and Privacy*. 179–186.

KABUS, P., TERPSTRA, W. W., CILIA, M., AND BUCHMANN, A. P. 2005. Addressing cheating in distributed MMOGs. In *Proceedings of 4th ACM SIGCOMM Workshop on Network and System Support for Games*.

KAISER, E., FENG, W., AND SCHLUESSLER, T. 2009. Fides: Remote anomaly-based cheat detection using client emulation. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*.

KRUEGEL, C., KIRDA, E., MUTZ, D., ROBERTSON, W., AND VIGNA, G. 2005. Automating mimicry attacks using static binary analysis. In *Proceedings of the 14th USENIX Security Symposium*. 161–176.

LYHYAOUI, Y., LYHYAOUI, A., AND NATKIN, S. 2005. Online games: Categorization of attacks. In *Proceedings of the International Conference on Computer as a Tool (EUROCON)*.

MAGIERA, M. 2009. Videogames sales bigger than DVD-Blu-ray for first time. http://www.videobusiness.com/article/CA6631456.html.

MITTERHOFER, S., PLATZER, C., KRUEGEL, C., AND KIRDA, E. 2009. Server-side bot detection in massive multiplayer online games. *IEEE Secu. Priv. 7,* 3, 18–25.

MÖNCH, C., GRIMEN, G., AND MIDTSTRAUM, R. 2006. Protecting online games against cheating. In *Proceedings of the 5th ACM SIGCOMM Workshop on Network and System Support for Games*.

MULLIGAN, J. AND PATROVSKY, B. 2003. *Developing Online Games: An Insider's Guide*. New Riders Publishing.

ROSENBLUM, M. AND OUSTERHOUT, J. K. 1992. The design and implementation of a log-structured file system. *ACM Trans. Comput. Syst. 10,* 1, 26–52.

SCHLUESSLER, T., GOGLIN, S., AND JOHNSON, E. 2007. Is a bot at the controls? Detecting input data attacks. In *Proceedings of the 6th ACM SIGCOMM Workshop on Network and System Support for Games*. 1–6.

SPOHN, D. Cheating in online games. http://internetgames.about.com/od/gamingnews/a/cheating.htm.

VIKRAM, K., PRATEEK, A., AND LIVSHITS, B. 2009. Ripley: Automatically securing Web 2.0 applications through replicated execution. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*.

WANG, R., WANG, X., LI, Z., TANG, H., REITER, M. K., AND DONG, Z. 2009. Privacy-preserving genomic computation through program specialization. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*.

WARD, M. 2005. Warcraft game maker in spying row. http://news.bbc.co.uk/2/hi/technology/4385050.stm.

WEBB, S. AND SOH, S. 2008. A survey on network game cheats and P2P solutions. *Aust. J. Intell. Inform. Process. Syst. 9,* 4, 34–43.

YAMPOLSKLY, R. V. AND GOVINDARAJU, V. 2007. Embedded noninteractive continuous bot detection. *Comput. Entertain. 5,* 4, 1–11.

YAN, J. AND RANDELL, B. 2005. A systematic classification of cheating in online games. In *Proceedings of the 4th ACM SIGCOMM Workshop on Network and System Support for Games*.

YANG, J., SAR, C., TWOHEY, P., CADAR, C., AND ENGLER, D. 2006. Automatically generating malicious disks using symbolic execution. In *Proceedings of the IEEE Symposium on Security and Privacy*.

Received July 2010; revised March 2011; accepted July 2011