


Physics Engine Development

Sergiy Migdalskiy
Dirk Gregorius



Physics
Rubikon



Hello!

My name is Sergiy Migdalskiy and I work at Valve with Dirk Gregorius. I helped **port**, develop and optimize our in-house physics engine, Rubikon. Also, I have implemented a physics engine for **Uncharted: Drake's Fortune** on PS/3.

Physics Engine Development



While this talk is about physics development, I'd like to emphasise that tools and techniques I'm describing are general purpose. They are probably applicable to any complex system written in C++.

Visualization



Run `Playing_with_Physics.mov`

This is what I do the whole day long.

I play with physics, throw ragdolls into sandboxes, and such silliness.

(wait) Here I'm piling up some ragdolls, it's looking good, until..

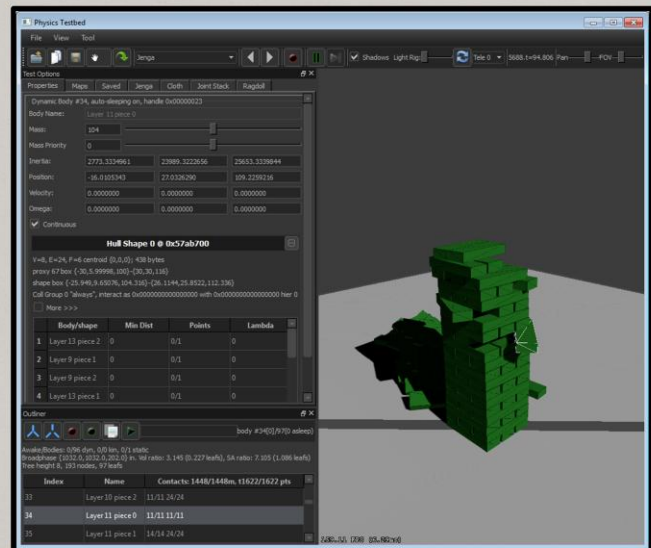

I notice one of the ragdolls, the very first, doesn't seem to collide with anything.

I'm testing some more. I make sure the other ragdolls work fine. Then I notice something.

Here it is – the collision flags on the very first ragdoll are set to not collide with anything.

Now the mystery is solved, and I go on to other tests.

Physics + Qt = Fast Dev Cycle



We started working with Dirk by porting his rigid body engine to Valve's **framework**. I knew Qt and I like playing with stuff I'm **working on**, for fun and to make it **easier**. So I implemented a little IDE for our physics **project**. It was incremental. Every time something took more than 30 seconds to look up in the debugger, I made a widget that makes it a split-second. We had physics running in an isolated tool with buttons, sliders and 3D window with **shadows**, the one you've just seen, which was pretty **cool**. And it restarted in **seconds**. We call it the Physics Testbed, or the Physics Debugger. Not to be confused with Visual Studio debugger.

Physics + Game = Slow Dev Cycle

Collision in Game



But there comes the time for every game physics engine to be put into an actual game.

Once we started integrating physics with the game, we had to debug it inside the game. So we lost our fast-iteration, nice UI.


Every engine has in-game visualization of physics. So we make do.

It has some serious drawbacks, though.

- You can't easily use mouse cursor for clicking collision or UI.
- You cannot fly around easily, you have to implement a special camera for it.
- You need to render inside game engine, which can be pretty tricky.
- You have to implement a special pause that pauses absolutely everything while you're examining something.
- It's non-trivial to implement rich UI, like you can do with Qt.

Physics UI

Collision in a Separate Window with UI

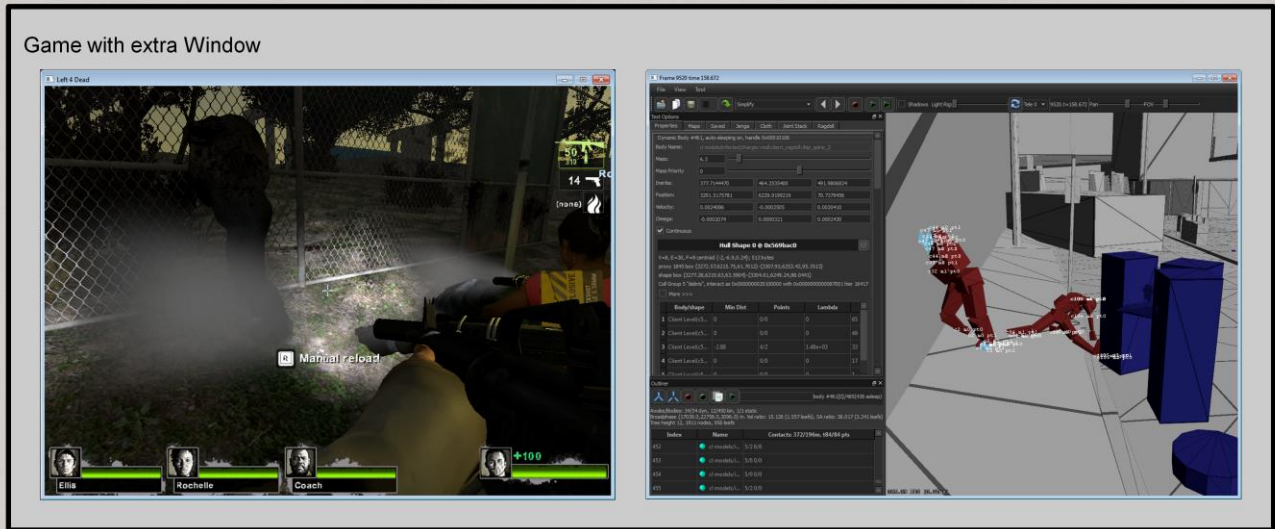


Physics Testbed was a great UI with a bunch physics-centric widgets. It was a great help already in making physics work. I felt I can reuse it in the game integrated physics.

But how do we display in-game physics world in a separate window?

Extra Window

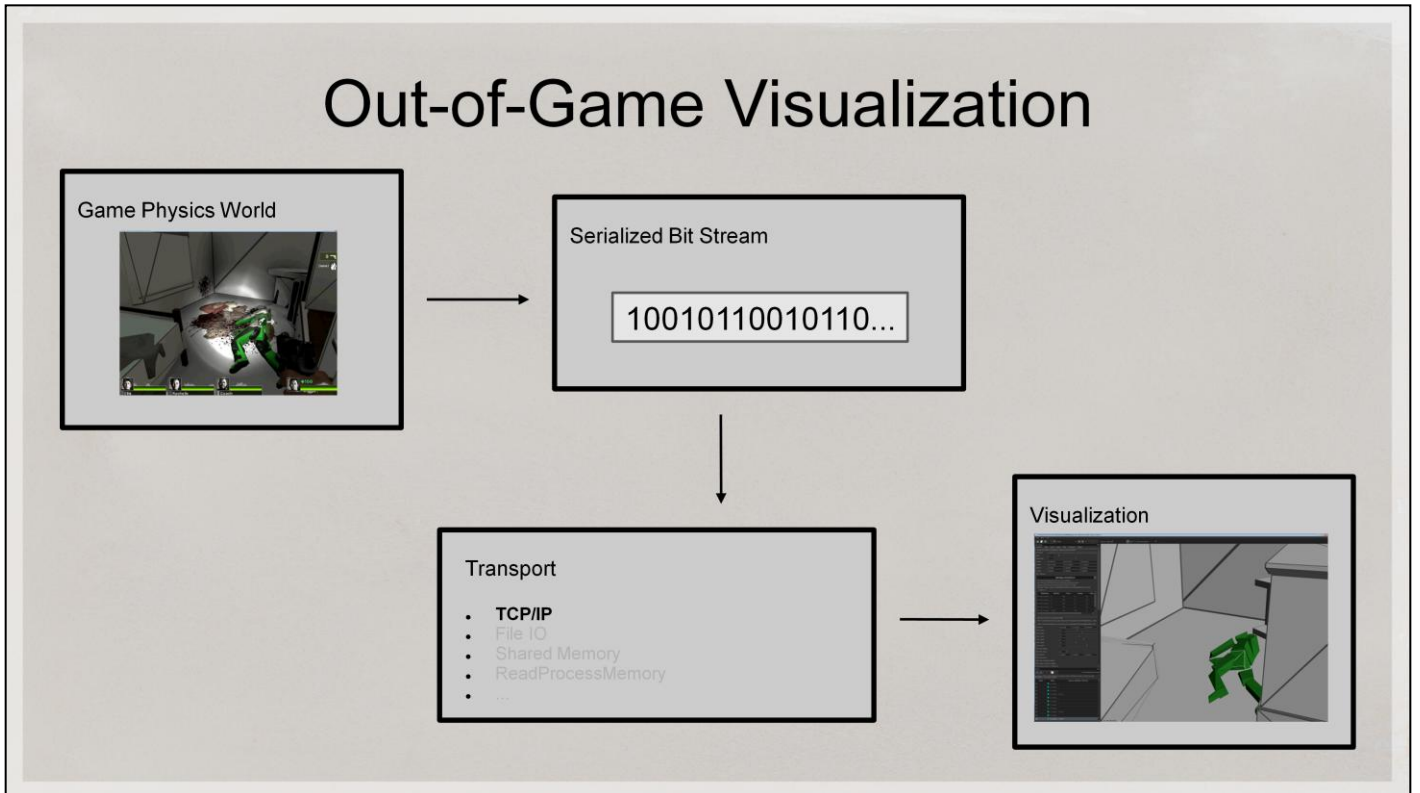
Game with extra Window



I briefly entertained an idea to open an extra Qt window in the game process. I find it:

- inconvenient to implement
 - It's intrusive to the main game loop
 - does not work when the game is crashed or stopped on a breakpoint
 - You have to implement "full pause" in the game, both on server and client
- So I decided against in-process window. This means I could just reuse the existing separate tool, our Physics Testbed.


Out-of-Game Visualization



To visualize in external process, you need to send data to it.
The hardest part about sending physics data to external app is serialization.
But if you can properly and fully serialize your world state into a stream of bits and send it over the wire, it becomes very straightforward to visualize your physics in a separate app.

Playing with physics in external app is a nice bonus feature.

Serialization

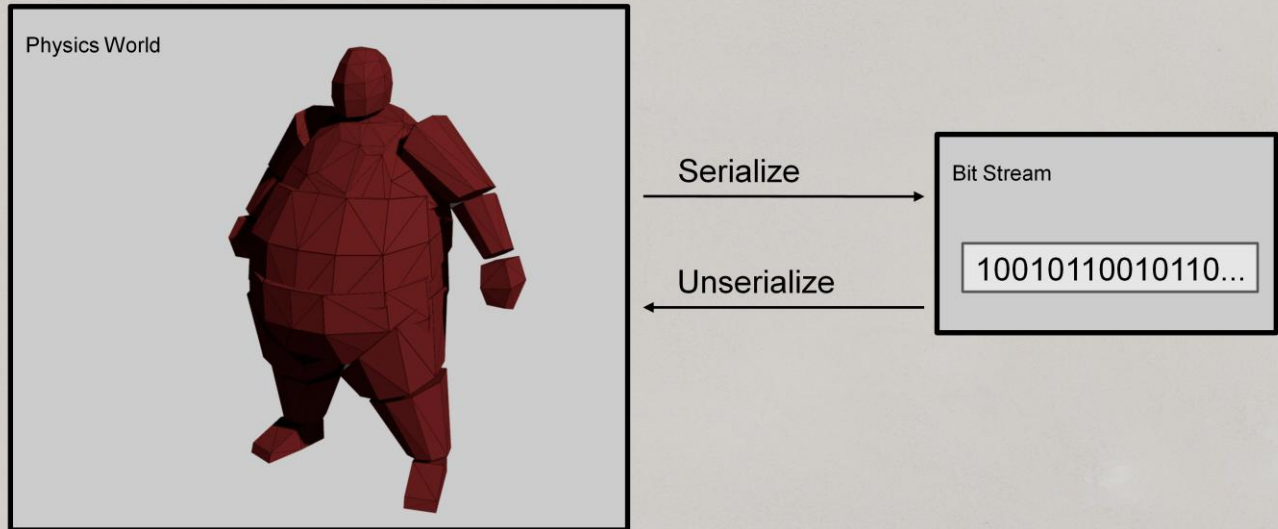


So serialization is the biggest obstacle here. But once we solve it, it becomes very useful:

- we can save and look at something later
- we can implement nice data and time profilers
- we can record and replay bugs

Once you can serialize, it is trivial to stream to external app.

(Un)Serialization

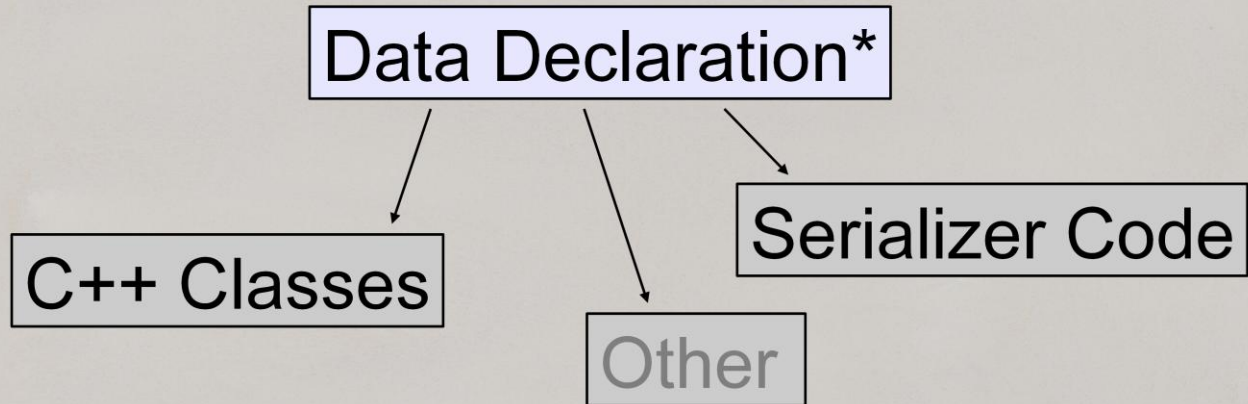


By the way, of course, serialization works two ways. The external app has to somehow turn our stream of bits into data structures..

So How do we serialize? Physics engine is complex. There are so many classes and data structures in there! Serializing pointers is hard. How do we build software that doesn't miss anything, and then reconstructs everything faithfully. And hopefully it updates itself every time we add or remove a field in a class?

Definitely not! by writing serialization by hand.

Generating C++

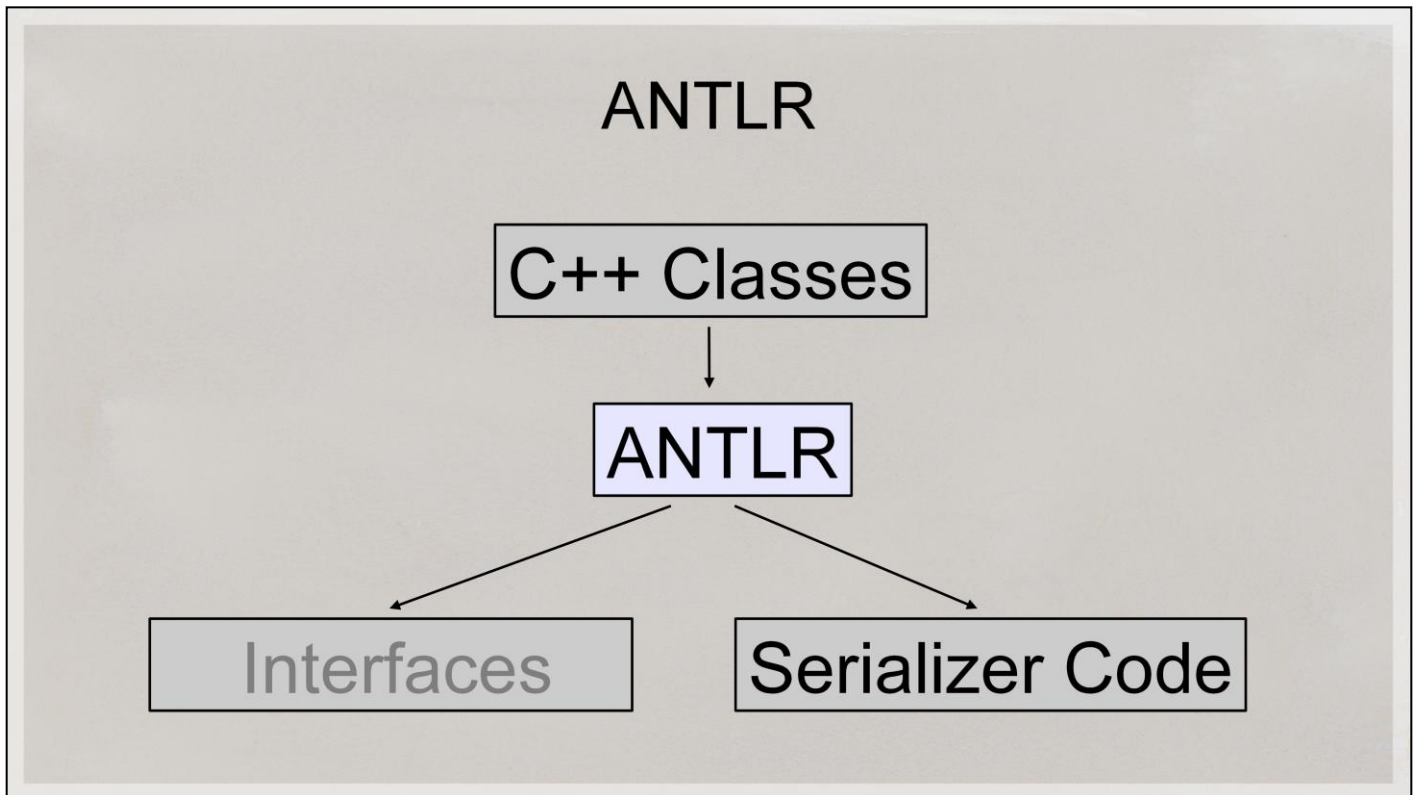


* Data Declarations in easy-to-parse language (not C++)

- One way to solve it is to describe your data structure in non-C++ (like protobufs), and generate C++ code from that.
- The description has to be easy to parse

We already had C++ code to start with, and the data structures were not easy to express in something like protobufs. I'd have to struggle to define the language and then struggle to express our nontrivial C++ data structures.

If only we use C++! Then there is no need to remember the new language syntax. Anybody can come in and modify the code without learning an extra language



- ANTLR is a parser generator and was the first thing I used. It has a whole IDE for debugging grammars. It's very powerful.
- In house, we use ANTLR for our reflection API. We only parse a subset of C++ specification. I took that parser and I still use it to generate clean public interfaces to our physics engine. They take care of type conversions, they log external activity and do some other things.
- ANTLR is not ideal for parsing C++, though. You have to maintain C++ grammar, Java is its native language so it's the easiest target to generate parsers. It's fast, but an optimized C++ parser is faster.

Generating C++ from C++

C++ Classes



Serializer Code

We can analyze the engine's header files to generate serialization code.

That's what we would do if we wrote the serialization by hand.

But instead of doing it by hand we can parse C++ in a utility and spit out C++ which will serialize our data.

But parsing C++ is extremely hard. Or is it?

Clang to the Rescue!

C++ Classes



Clang




Serializer Code

- Clang is an open-source c++ compiler. In some ways it's a direct competitor of GCC, but with much more readable source code and liberal license.
- Clang makes C++ parsing a piece of cake. It's written in C++ itself, and it just works.
- We can use all our existing code as input, convert cpp and header files into easy-to-use AST and spit out the serializer routines.
- One bonus from using Clang is this improved C++ compliance of our code. We found some bugs, and prepared the code for Linux. It's really useful to compile your code with another compiler.

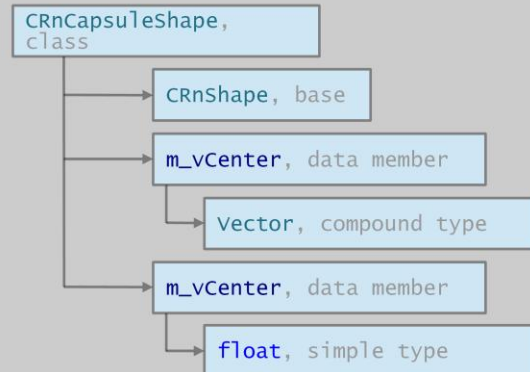
Clang AST

Source code

```
class CRnCapsuleShape :  
    public CRnShape  
{  
    public:  
    ...  
    private:  
        Vector m_vCenter[ 2 ];  
        float m_flRadius;  
        AUTO_SERIALIZE;  
};
```



Abstract Syntax Tree



I mentioned the term AST. It means Abstract Syntax Tree.

It's a tree that describes everything there is to know about your source code.

This is how it looks, schematically. For every class, you'll have a node. For every member, its type or size, or annotation, or anything you can think of you'll probably get a node of some type.

Code Generation



```
#include "..."  
class CRnCapsuleShape :  
    public CRnShape  
{  
public:  
    ...  
private:  
    Vector m_vCenter[ 2 ];  
    float m_fRadius;  
    AUTO_SERIALIZE;  
};
```

```
void CRnCapsuleShape::Serialize(  
    CRnSerializer *pOut ) const  
{  
    CRnShape::Serialize( pOut );  
    pOut->WriteBuiltin<float>( m_fRadius );  
    for( int nElement = 0;  
        nElement < 2;  
        nElement )  
    {  
        ::Serialize( pOut,  
                    m_vCenter[nElement] );  
    }  
}
```

- In effect, we are transforming C++ code into additional C++ code
- It's very tempting at first to put code generation as a bunch of printf() statements right into your parser
- I decided to resist the urge, and I consider myself lucky I did. It would be horribly slow to iterate on the generated code with printf's.

The Structured Approach



```
#include "..."  
class CRnCapsuleShape :  
    public CRnShape  
{  
    public:  
    ...  
    private:  
        Vector m_vCenter[ 2 ];  
        float m_fIRadius;  
        AUTO_SERIALIZE;  
};
```


```
"CRnCapsuleShape" : {  
    "fields" : {  
        "m_fIRadius" : {  
            "typeName" : "float",  
        },  
        "m_vCenter" : {  
            "className" : "vector",  
            "arraySize" : 2,  
        }  
    },  
    "bases" : [  
        "CRnShape"  
    ],  
},
```

Parse, *then* Generate. Decouple these two stages, and unless your project is trivial, you'll be rewarded.

I'm distilling the huge AST of the physics engine into a very concise json file. Json is a human readable format that is widely used nowadays.


After adding features for a year, my Visitor class in the code that uses Clang that just prints out the json with relevant information is about 1000 lines long, and it would be incomprehensible if I generated code there.

Json?



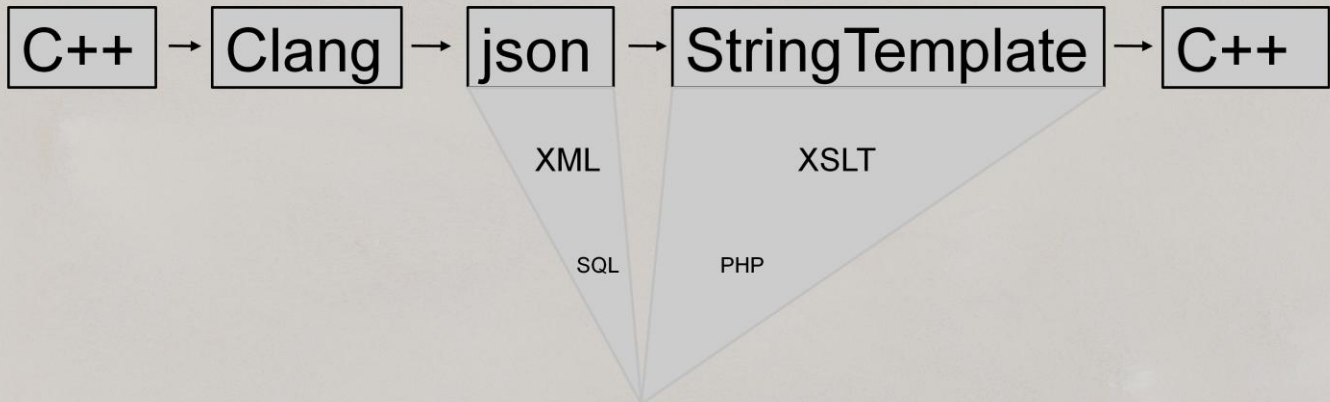
- So, I'm taking my C++ code, let Clang parse it and spit out Json.
- **But I ultimately need my serializer and potentially other useful tools to be written in C++**
- This begs the question: how do I generate C++ code from Json?

StringTemplate (to the rescue)!




- Fortunately, there is no need to reinvent the wheel, because it was invented many times over before us. There are a lot of text template libraries, and languages. They are widely used and freely available.
- I chose StringTemplate. It's a text library that's very easy to use, and it's very well suited to code generation. It's declarative style, which is appealing for this kind of task.

StringTemplate (to the rescue)!



- You could use something else. In fact, if you're very familiar and comfortable with, say, PHP, you should absolutely use it.
- You could also use XML instead of Json. In fact, standard Clang tools can output full abstract syntax tree in XML directly. XSLT is a powerful transformation language for XML. StringTemplates and Json are much easier to learn and much more human-readable. That's why I personally chose them.
- There are many ways to configure this pipeline, I'm just describing what worked for me.


StringTemplate



```
void <name>::Serialize(  
    CRnSerializer *pout ) const  
{  
    <class.bases :  
        {b | <b>::Serialize( pout ); }>  
    <class.fields.keys, class.fields.values: {k,  
v | <field( name = k, props = v )> }>  
}
```

- StringTemplate looks like this. The tool itself is written in java, but there's no need to know Java to use it. All my serialization, unserialization and memory statistics code is generated from a 400-line template. These are a few of those lines. They auto-generate the code to serialize most classes in our physics engine.
- You feed it json file, it spits out C++ file. Quite simple.

Code Generation: Sources



```
#include "..."  
class CRnCapsuleShape :  
    public CRnShape  
{  
    public:  
    ...  
    private:  
        Vector m_vCenter[ 2 ];  
        float m_fIRadius;  
        AUTO_SERIALIZE;  
};
```

- So, to recap.
- We start with the source code of the engine – it's arbitrary C++, potentially annotated. For instance, you can annotate only the classes you want serialized. And exclude some fields from serialization. And add a member to call after serialization.


Code Generation: Parser



```
clang::CXXMethodDecl * FindMethodswithName( const clang::CXXRecordDecl *pRecord,
                                             const char *pName )
{
    for( auto itMethod = pRecord->method_begin(),
          itMethodEnd = pRecord->method_end();
          itMethod != itMethodEnd; ++itMethod )
    {
        if( clang::IdentifierInfo *pIdInfo = ( *itMethod )->getIdentifier() )
        {
            if( pIdInfo->getName() == pName )
            {
                return *itMethod;
            }
        }
    }
    return NULL;
}
```

- The parser/Clang part is moderately complicated. I wrote the first working version within a couple of weeks, working on it part-time. But then I fiddled with it for well over a year. It handles all the special and edge cases we have in our engine. It's not a magic bullet, though. It will not parse anything and everything in the known universe, but it'll parse and describe any exotic data structure in our engine. I have to change it only very rarely. In fact, the last change I made was very minor. And it was over 4 months ago.
- This slide shows how to look for a method with a given name in class. It's pretty straightforward to extract information once you start writing the code.
- The best source of information for me was the online Clang docs, especially their class diagram. It's very logical and educational. Also, Clang is very readable, so a lot of stuff becomes clearer after you debug it.

Code Generation: Parser



```
"CRnCapsuleShape" : {  
  "fields" : {  
    "m_fRadius" : {  
      "typeName" : "float",  
    },  
    "m_vCenter" : {  
      "className" : "Vector",  
      "arraySize" : 2,  
    }  
  },  
  "bases" : [  
    "CRnShape"  
  ]  
},  
...
```

- The generated json file looks like this. I already explained the meaning of this snippet before.
- It contains all the data necessary to generate the serializer and whatever other tool you might wish to generate, and nothing else. It's very much human-readable, which helps a lot when developing it
- On the slide, it shows the class name, the list of its fields with field names and types and if the field is an array, then the array count. There's also the base class name. It's what you would expect.

Generated Code

C++

Clang

json

StringTemplate

C++

```
void <name>::Serialize(  
    CRnSerializer *pOut ) const  
{  
    <class.bases :  
        {b | <b>::Serialize( pOut ); }>  
    <class.fields.keys, class.fields.values:  
{k, v | <field( name = k, props = v )> }>  
}
```

```
void CRnCapsuleShape::Serialize(  
    CRnSerializer *pOut ) const  
{  
    CRnShape::Serialize( pOut );  
    pOut->WriteBuiltin<float>( m_fRadius );  
    for( int nElement = 0;  
        nElement < 2;  
        nElement )  
    {  
        ::Serialize( pOut,  
                    m_vCenter[nElement] );  
    }  
}
```

- To remind, this is what the string template looks like, and this here is the resulting generated C++ code
- The template goes through all data members of the class and generates code for each of them.
- The template is smart enough to know where I have arrays, templates, pointers to follow and so on.
- All the non-plain data elements are recursively serialized by other overloaded serialization functions. Some of them methods, some of them free functions. I especially prefer free functions for simpler types, containers and third party classes where I don't want to touch the header where those classes are declared.



Once you complete these 5 easy steps, you'll have yourself one wholesome serializer. It will be self-maintaining. You will forever be able to visualize your physics in all detail. Sometimes we don't even know something is wrong. Some bugs ship and go unnoticed for years. This will never happen to you again. Well..., I can't make guarantees... But I really hope I convinced you that visualizing stuff is great. Having serializer auto-magically generate itself makes everything even better. But we have quite a bit more to discuss before we wrap up.

Record and Replay



=Play= Top_Heavy_Lamp.wmv

Here's an example of one workflow when you have full serialization.
I noticed something weird with a lamp here. It was always falling down. Lamps are supposed to stand straight, and this one was always on its side.
So I snooped and looked at the lamp.
(wait) When I figured out why it's not standing (because its center of mass is too high), I just sent out an email to an artists – and I attached the file he can play with to prove my words. It's quite intuitive.

Clang's “Hello World!”

```
See ClangCheck.cpp

class CDumpAction: public clang::ASTConsumer
{
    llvm::raw_fd_ostream &m_log;
public:
    CDumpAction( llvm::raw_fd_ostream &fd ):m_log( fd )
    {
    }


    virtual bool HandleTopLevelDecl( clang::DeclGroupRef DG )
    {
        for ( clang::DeclGroupRef::iterator it = DG.begin(),
              itEnd = DG.end(); it != itEnd; ++it )
        {
            clang::Decl *pDecl = *it;
            pDecl->dumpXML( m_log );
        }
        return true;
    }
};
```

Here is the Clang's “Hello World” program relevant to this use case is called ClangCheck and it is a great starting point to write your own tool. It's part of the Clang distro, and I urge you to have a look at it.

It uses visitor pattern. For every top-level declaration in your source file, such as a class or function definition, or a typedef, or a global, you get a callback. In that callback, you can drill down the AST.

This self-contained example is dumping those declarations as XML to a log. If XML is your thing and performance is not the most critical issue, this here just may be your parser.

What to Parse?



Typically, all our class declarations are contained in headers.

It's enough to parse one .cpp file that #includes all the relevant headers for serialization. There's no need to parse the whole engine, although there may be other useful tools that would need it.

To parse just one file, you'll need to create a compilation database file with cpp file name and compile options.

Clang Compilation Database

```
[{  
  "directory": "c:\\foo\\blah",  
  "command": "clang -fsyntax-only -fms-extensions -fms-compatibility -I../public/include .....  
              rnserialize.cpp",  
  "file": "rnserialize.cpp"  
}]
```

Here's an example of compilation database. The most interesting options are highlighted here.

Compilation database is like a makefile, or vcproj. If you don't want to do anything non-standard, you can simply generate a .json file with all the information and give it to clang. We use a tool to generate VCPROJ files, and we changed that tool to also spit out compilation databases.

Clang Bootstrap Cheatsheet

```
class MyAction: public clang::ASTConsumer
{
public:
    virtual bool HandleTopLevelDecl ( clang::DeclGroupRef DG )
    {
        for ( clang::DeclGroupRef::iterator
              it = DG.begin(), itEnd = DG.end(); it != itEnd; ++it )
        {
            clang::CXXRecordDecl *pRecord =
                llvm::dyn_cast<clang::CXXRecordDecl>( *it );
            // do something with it...
        }
        return true;
    }
};
```

```
class MyFactory
{
public:
    clang::ASTConsumer
        *newASTConsumer()
    {
        return
            new MyAction( );
    }
};
```


```
pCompilationDb = clang::tooling::JSONCompilationDatabase::loadFromFile( strJsonDb, errorMessage );
files = pJsonDb->getAllFiles();
clang::tooling::ClangTool Tool( *pCompilationDb, files );
MyFactory factory( pJsonDbPath );
int nError = Tool.run( clang::tooling::newFrontendActionFactory( &factory ) );
llvm::errs().flush();
llvm::outs().flush();
```

And here is a little cheatsheet. In your parser tool, that you carefully link with all the clang libraries, what do you do first?

Start with implementing ASTConsumer, make a factory class for it, and then in your main() function, load compilation database and then run an instance ClangTool. It's all here in this slide.


One thing to remember: you may have to flush std out and std error.

Box Cast



Before going further, I'd like to clarify what I mean by box cast. We often use it in Source engine. It's like raycast, but with a box.

Contact



Just take a box, move it linearly until you hit something.
At that point, you sometimes create a contact point for the solver.
We use it for player movement a lot, so debugging it is a common thing to do.

Visualizing Physics In Game



=Run= Player_Movement.mov

Let's see what we've got for our efforts so far.

Here's a game running alongside visualization tool.

I'm walking around a room, and I'm looking at the raycasts and box casts our player movement is making

I can look at it from the 3rd person view, I can only render the relevant data.

I'm noticing something strange, there's a flurry of box casts sometimes. It might probably result in a framerate hitch.


So I'm trying to reproduce this.

I find a place where it's happening.

And I'm flying around to see what causes it.

So far, so good. By the way, I can play with that physics at any point, but it won't affect the game, so it's not useful in this case.

When That Is Not Enough




Well, what if just looking at the problem does not present an obvious fix?

What if we have to actually debug code?

Debugging


Engine Parts



At least when you debug your own code, it's easy to see the parts. The components.
What they are supposed to do.
It's easy to check and see if they are actually doing it.

Other People's Code

Unfamiliar Code Looks Like This:



Working in a team, it's different... It's hard to see the whole picture. It's hard to check if the parts are working properly.

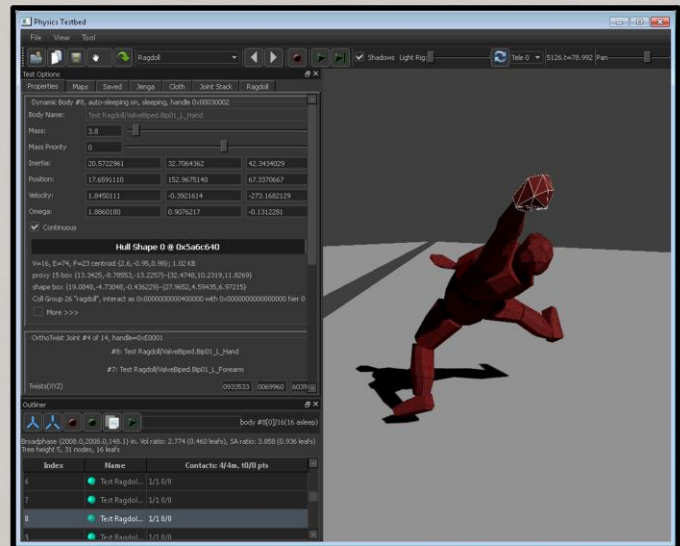

Other People's Bugs

Bug in Unfamiliar Code Is Here:



Even when I find a bug in foreign code, there are usually many ways to fix it and I'm never quite sure which one is the best.
Or maybe I misunderstood the code and my fix will cause another bug.

Second Life of Testbed




And now we have this wonderful application.. I wonder if we could use it when we step through the breakpoints just like we use it when we walk through the game. We need to leverage it! But how?

Debug Experience

```
...
 * m_pParticles,14 0x0000000011266300 pos=(- 6876.29443, 6769.97656, 364.533
 * [0] pos=(- 6876.29443, 6769.97656, 364.533264) x=( 0.231891751,
 * [1] pos=(- 6854.43750, 6699.25000, 336.452576) x=(0.231891751,
 * [2] pos=(- 6866.72510, 6739.01172, 352.239136) x=(0.231891751,
 * [3] pos=(- 6896.30566, 6721.00684, 412.660950) x=( 0.621897459,
 * [4] pos=(- 6903.90430, 6718.92432, 419.238647) x=( 0.795030117,
 * [5] pos=(- 6892.56396, 6721.99951, 408.815186) x=( 0.576003432,
 * [6] pos=(- 6877.71094, 6774.56104, 366.353394) x=( 0.231891751,
 * [7] pos=(- 6865.03955, 6733.55664, 350.073364) x=(0.231891751,
 * [8] pos=(- 6915.52539, 6715.27734, 420.790619) x=( 0.745445311,
 * [9] pos=(- 6862.48145, 6725.27881, 346.786865) x=(0.231891751,
 * [10] pos=(- 6925.21631, 6711.95996, 417.005035) x=( 0.674079120,
 * [11] pos=(- 6882.99365, 6781.65576, 373.140333) x=( 0.231891751,
 * [12] pos=(- 6887.22559, 6805.36816, 378.716858) x=( 0.231891751,
 * [13] pos=(- 6899.58838, 6845.37256, 394.600006) x=( 0.231891751,
 * pSimulationWorldTransforms,14 0x00000000579603a0 (pos=( 6904.11133, 6710.47217, 325.63
 * [0] pos=( 6904.11133, 6710.47217, 325.638428) x=(0.772891462,
 * * matrix3x4_t 0x00000000579603a0 (0.772891462, 0.3
 * * m_fMatVal 0x00000000579603a0 (0.0021016858
 * [1] pos=( 6901.78857, 6711.84180, 349.343170) x=( 0.021016858
 * [2] pos=( 6896.30271, 6697.04053, 322.380066) x=(0.372540712,
 * [3] pos=( 6907.72754, 6725.45947, 321.857361) x=( 0.086953304,
 * [4] pos=( 6902.04639, 6711.83691, 364.934204) x=(0.0402162559,
 * [5] pos=( 6900.82715, 6712.28955, 380.599396) x=(0.0578526743,
 * [6] pos=( 6900.03711, 6714.20801, 391.881500) x=(0.610547960,
 * [7] pos=( 6892.98975, 6711.08350, 381.304016) x=(0.299724787,
 * [8] pos=( 6899.02979, 6720.03516, 380.205017) x=(0.599534094,
 * [9] pos=( 6890.17285, 6721.56641, 397.602539) x=(0.798212171,
 * [10] pos=( 6879.02051, 6716.74902, 393.418640) x=(0.136251017,
 * [11] pos=( 6883.49463, 6730.69482, 392.627655) x=(0.263435274,
 * [12] pos=( 6892.62842, 6721.41406, 408.810089) x=( 0.579993010,
 * [13] pos=( 6884.52002, 6723.33447, 396.799957) x=(0.256185055,
 * * m_DebugName (m_pName=0x0000000000000000 <NULL> m_DebugIndex=0
 * * m_pWorld 0x0000000020bf700 (m_Filter=m_GroupPairs=0x0000000002
 * * m_pFeModel 0x00000000110384f9 (m_nFlags=4009750271 m_nNodeCount
 * * m_fBendUnderRelax 0.000000000
 * * m_fStretchUnderRelax 0.000000000
 * * m_fOverPredict 0.000000000
 * * m_nNodeCount 14
 * * m_nParticleCount 14
 * * m_pParticles 0x0000000011266300 pos=(- 6876.29443, 6769.97656, 364.533
 * * m_pPos0 0x0000000011266840 -1.#IND00000, -1.#IND00000, -1.#IND00000
 * * m_pPos1 0x0000000011266920 -1.#IND00000, -1.#IND00000, -1.#IND00000
 * * m_nSimFlags 4
 * * m_nSimFlags 4

```



```
...
 * m_pParticles,14 0x0000000011266300 pos=(- 6876.29443, 6769.97656, 364.53
 * [0] pos=(- 6876.29443, 6769.97656, 364.533264) x=( 0.231891751,
 * [1] pos=(- 6854.43750, 6699.25000, 336.452576) x=(0.231891751,
 * [2] pos=(- 6866.72510, 6739.01172, 352.239136) x=(0.231891751,
 * [3] pos=(- 6896.30566, 6721.00684, 412.660950) x=( 0.621897459,
 * [4] pos=(- 6903.90430, 6718.92432, 419.238647) x=( 0.795030117,
 * [5] pos=(- 6892.56396, 6721.99951, 408.815186) x=( 0.576003432,
 * [6] pos=(- 6877.71094, 6774.56104, 366.353394) x=( 0.231891751,
 * [7] pos=(- 6865.03955, 6733.55664, 350.073364) x=(0.231891751,
 * [8] pos=(- 6915.52539, 6715.27734, 420.790619) x=( 0.745445311,
 * [9] pos=(- 6862.48145, 6725.27881, 346.786865) x=(0.231891751,
 * [10] pos=(- 6925.21631, 6711.95996, 417.005035) x=( 0.674079120,
 * [11] pos=(- 6882.99365, 6781.65576, 373.140333) x=( 0.231891751,
 * [12] pos=(- 6887.22559, 6805.36816, 378.716858) x=( 0.231891751,
 * [13] pos=(- 6899.58838, 6845.37256, 394.600006) x=( 0.231891751,
 * pSimulationWorldTransforms,14 0x00000000579603a0 (pos=( 6904.11133, 6710.47217, 325.63
 * [0] pos=( 6904.11133, 6710.47217, 325.638428) x=(0.772891462,
 * * matrix3x4_t 0x00000000579603a0 (0.772891462, 0.3
 * * m_fMatVal 0x00000000579603a0 (0.0021016858
 * [1] pos=( 6901.78857, 6711.84180, 349.343170) x=( 0.021016858
 * [2] pos=( 6896.30271, 6697.04053, 322.380066) x=(0.372540712,
 * [3] pos=( 6907.72754, 6725.45947, 321.857361) x=( 0.086953304,
 * [4] pos=( 6902.04639, 6711.83691, 364.934204) x=(0.0402162559,
 * [5] pos=( 6900.82715, 6712.28955, 380.599396) x=(0.0578526743,
 * [6] pos=( 6900.03711, 6714.20801, 391.881500) x=(0.610547960,
 * [7] pos=( 6892.98975, 6711.08350, 381.304016) x=(0.299724787,
 * [8] pos=( 6899.02979, 6720.03516, 380.205017) x=(0.599534094,
 * [9] pos=( 6890.17285, 6721.56641, 397.602539) x=(0.798212171,
 * [10] pos=( 6879.02051, 6716.74902, 393.418640) x=(0.136251017,
 * [11] pos=( 6883.49463, 6730.69482, 392.627655) x=(0.263435274,
 * [12] pos=( 6892.62842, 6721.41406, 408.810089) x=( 0.579993010,
 * [13] pos=( 6884.52002, 6723.33447, 396.799957) x=(0.256185055,
 * * m_DebugName (m_pName=0x0000000000000000 <NULL> m_DebugIndex=0
 * * m_pWorld 0x0000000020bf700 (m_Filter=m_GroupPairs=0x0000000002
 * * m_pFeModel 0x00000000110384f9 (m_nFlags=4009750271 m_nNodeCount
 * * m_fBendUnderRelax 0.000000000
 * * m_fStretchUnderRelax 0.000000000
 * * m_fOverPredict 0.000000000
 * * m_nNodeCount 14
 * * m_nParticleCount 14
 * * m_pParticles 0x0000000011266300 pos=(- 6876.29443, 6769.97656, 364.533
 * * m_pPos0 0x0000000011266840 -1.#IND00000, -1.#IND00000, -1.#IND00000
 * * m_pPos1 0x0000000011266920 -1.#IND00000, -1.#IND00000, -1.#IND00000
 * * m_nSimFlags 4
 * * m_nSimFlags 4

```

The usual debug experience consists of sifting through Watch windows, stepping through lines of code, imagining our 3D world in our mind's eye..

Improved Debug Experience

```
0001
...
* m_Particles,14 0x0000000011266300 pos=( 6876.29443, 6769.97656, 364.533
+ [0] pos=( 6876.29443, 6769.97656, 364.533264) x=( 0.231891751,
+ [1] pos=( 6854.43750, 6699.25000, 336.452576) x=(0.231891751,
+ [2] pos=( 6866.72510, 6739.01172, 352.239136) x=(0.231891751,
+ [3] pos=( 6896.30966, 6721.00844, 412.660950) x=( 0.231891751,
+ [4] pos=( 6903.90430, 6718.92432, 419.238674) x=( 0.795030117,
+ [5] pos=( 6892.56396, 6721.99951, 408.815186) x=( 0.576003432,
+ [6] pos=( 6877.71094, 6774.56104, 366.353394) x=( 0.231891751,
+ [7] pos=( 6865.03955, 6733.55664, 350.073364) x=(0.231891751,
+ [8] pos=( 6915.52539, 6715.27734, 420.790619) x=( 0.745445311,
+ [9] pos=( 6862.48145, 6725.27881, 346.786865) x=(0.231891751,
+ [10] pos=( 6925.21631, 6711.95996, 417.005035) x=( 0.674079120,
+ [11] pos=( 6882.99365, 6781.65576, 372.140333) x=( 0.231891751,
+ [12] pos=( 6887.22559, 6805.36816, 378.716858) x=( 0.231891751,
+ [13] pos=( 6899.58838, 6845.37256, 394.600006) x=( 0.231891751,
* pSimulationWorldTransforms,14 0x0000000579603a0 ipos=( 6904.11133, 6710.47217, 325.63
+ [0] pos=( 6904.11133, 6710.47217, 325.638428) x=(0.772891462,
+ [1] * m_Matrix3x4_t 0x0000000579603a0 (0.772891462, 0.8
+ [2] pos=( 6901.78857, 6711.84180, 349.343170) x=( 0.021016858,
+ [3] pos=( 6896.30278, 6697.04029, 322.380966) x=(0.372540712,
+ [4] pos=( 6907.72754, 6725.45947, 321.857361) x=( 0.086953304,
+ [5] pos=( 6902.04639, 6711.83691, 364.934204) x=(0.0402162539,
+ [6] pos=( 6900.82715, 6712.28955, 380.599396) x=(0.0578526743,
+ [7] pos=( 6900.03711, 6714.20801, 391.881500) x=(0.610547968,
+ [8] pos=( 6892.98975, 6711.08350, 381.304016) x=(0.299724787,
+ [9] pos=( 6899.02979, 6720.03516, 380.205017) x=(0.599534094,
+ [10] pos=( 6896.17285, 6721.56641, 397.602339) x=(0.798212171,
+ [11] pos=( 6879.02051, 6716.74902, 393.418640) x=(0.136251017,
+ [12] pos=( 6883.49463, 6730.69482, 392.627655) x=(0.263435274,
+ [13] pos=( 6892.62842, 6721.41406, 408.810089) x=( 0.579993010,
+ [14] pos=( 6884.52002, 6723.33447, 396.799957) x=(0.236185053,
...
* m_DebugName 0x00000000110384f9 ipName=0x0000000000000000 -NULL- m_nDebugIndex=1
* m_pWorld 0x000000002bdf700 ipFilter=(m_GroupPairs=0x0000000000000000
* m_pFeModel 0x00000000110384f9 ipFilter=(m_GroupPairs=0x0000000000000000
* m_pBendUnderRelax 0.000000000
* m_pStretchUnderRelax 0.000000000
* m_pCoverPredict 0.000000000
* m_nNodeCount 14
* m_nParticleCount 14
* m_pParticles 0x0000000011266300 pos=( 6876.29443, 6769.97656, 364.533
* m_pPos0 0x0000000011266840 -1.#IND00000, -1.#IND00000, -1.#IND00000
* m_pPos1 0x0000000011266920 -1.#IND00000, -1.#IND00000, -1.#IND00000
* m_nSimFlags 4
...
m_SimTime 0.000000000
```



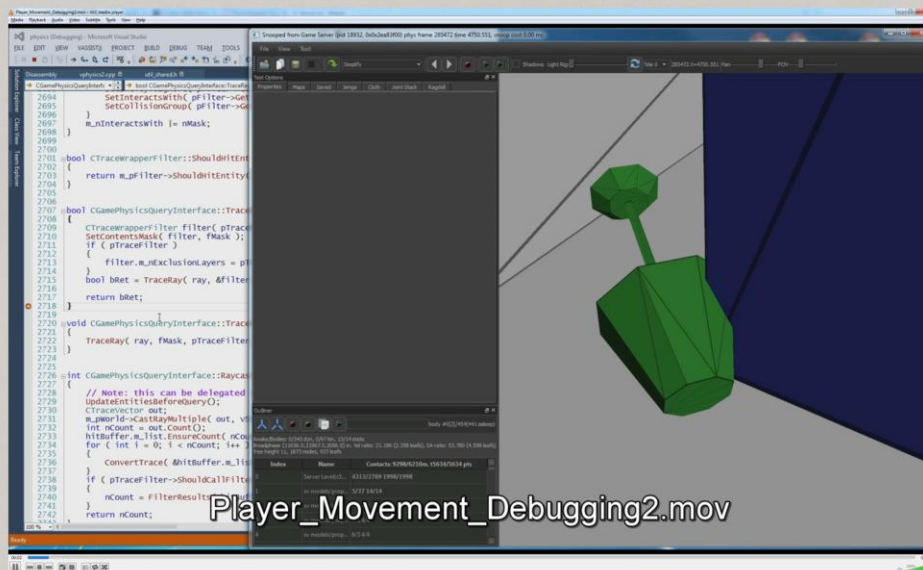
Ideally, I want to see the world just like I did in the Physics Testbed, but I also want to see it changing just in time as I step through my breakpoints.

Ideally, I want to see the world just like I did in the Physics Testbed, but I also want to see it changing just in time as I step through my breakpoints.

I want to be able to fly around the physics world frozen at a breakpoint. That would make debugging much more comfortable, enjoyable and productive.

That would make debugging much more comfortable, enjoyable and productive.

Debugging Visually



Let me show you what I mean.

As you can see, this is the place with too many box casts that I discovered in the previous demo.

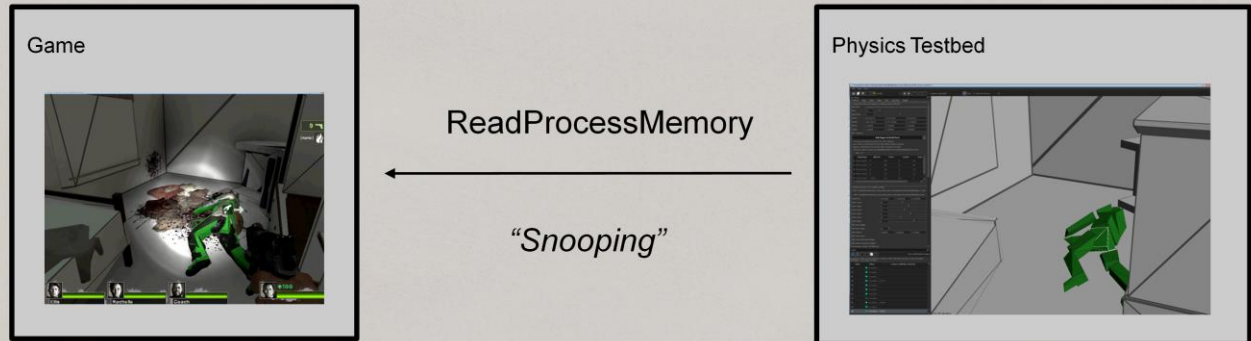
I place a breakpoint in the box trace function. Every time the it hits, I'm looking at the 3D world. I'm seeing the very last box cast.

This lets me step through bunch of casts, visualizing each of them in Physics Testbed.

This whole thing happens inside of a frame.

Please note that the game didn't send any data packets to Physics Testbed. It's casting boxes in a very tight loop. The testbed is effectively acting like a watch window.

Reading Memory Directly



The method that lets us do it is ReadProcessMemory API.

The Physics Testbed, acting in a way like a debugger, reads memory from the game without the game knowing it. I call this snooping. There is no way for the debugger to alter the game state, which is a nice benefit.

It works when process is stopped in debugger


There is no activity required from the game side besides a mutex and advertising where to start the search (the root data structure). Both are trivial to implement and are not intrusive.

ReadProcessMemory performance



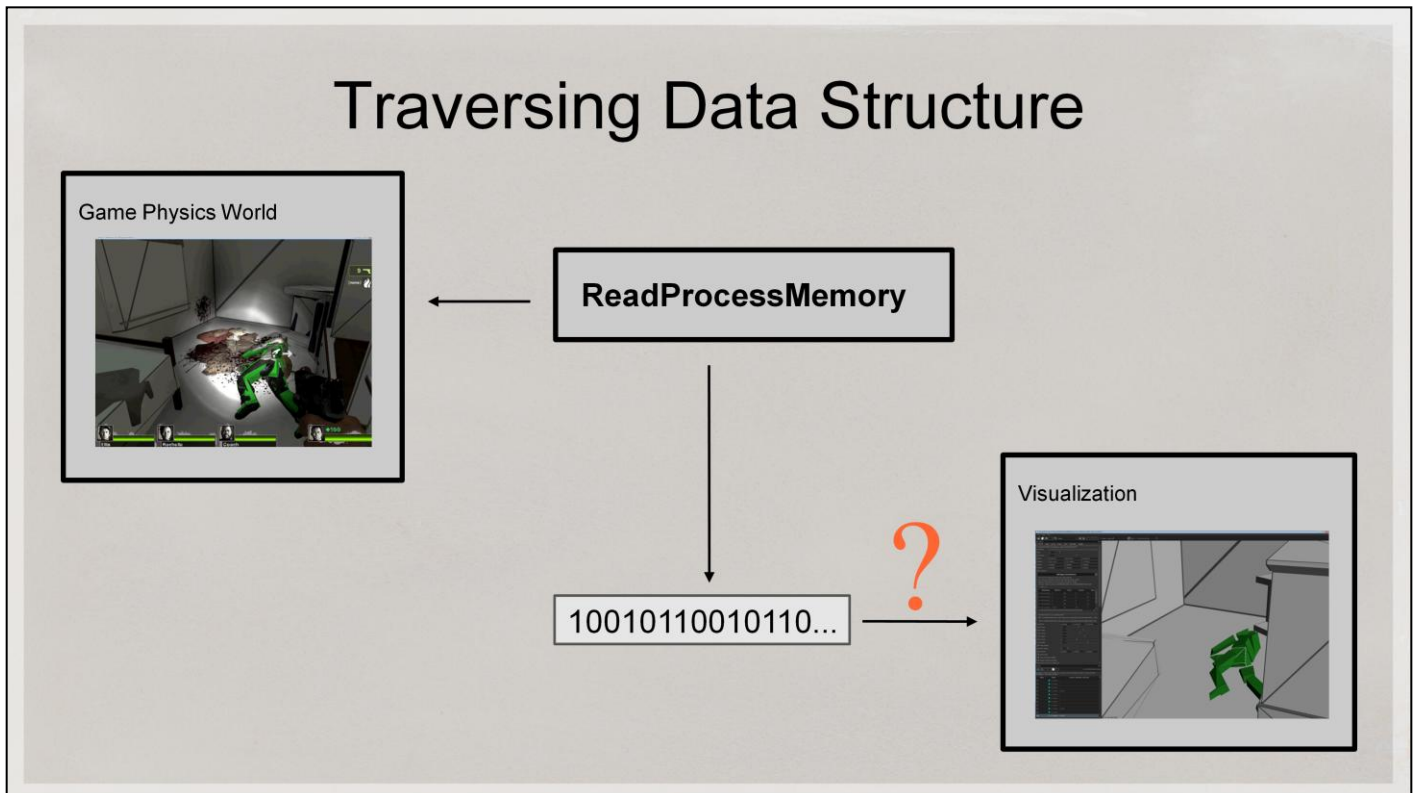
- The biggest question when I thought of this idea was whether it's fast enough. So I wrote a small benchmark.
- It is slow to read every 4- or 8-byte field across the Process boundary
- Every call to ReadProcessMemory descends into kernel, performs a syscall — it costs you at least 1us, limiting your bandwidth.
- But just copying 4 bytes in-process takes a thousandth of that time.

ReadProcessMemory performance



- I simply implemented software cache layer that reads at least 4Kb page every time you snoop a byte and caches that page
- You always know you can read the 4KiB page if you can read a single byte of it. There will be no memory protection faults.
- This cache helps a lot with reading C NULL-terminated strings, for which length is unknown until you read the whole string one byte at a time.

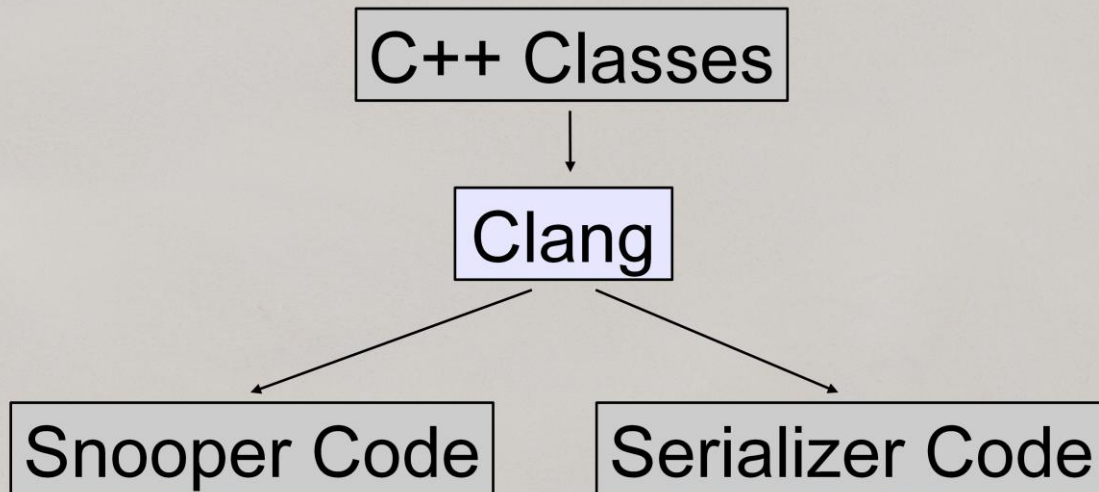
Traversing Data Structure



So, we can snoop the bytes from the game physics, but how do we know what to do with them?

In fact, walking the memory of the other process is very much like serialization. If you know your data structures, you can traverse them in another process using `ReadProcessMemory` just the same as you traverse them in the same process when you serialize them.

Snooper and Serializer



- We can still use Clang to understand our data structures.
- We can use all our existing code as input and spit out the snooper routines.

Snooper: Code Generation



```
#include "..."  
class CRnCapsuleShape :  
    public CRnShape  
{  
    public:  
    ...  
    private:  
        Vector m_vCenter[ 2 ];  
        float m_fRadius;  
        AUTO_SERIALIZE;  
};
```

```
void CRnCapsuleShape::Snoop( CRnSnooper*pIn,  
                             const CRnCapsuleShape *pLocalCopy )  
{  
    CRnShape::Snoop( pIn, pLocalCopy );  
    m_fRadius = pLocalCopy->m_fRadius;  
    for( int nElement = 0;  
         nElement < 2; ++nElement )  
    {  
        ::Snoop( pIn,  
                 &pLocalCopy->m_vCenter[nElement],  
                 m_vCenter[nElement] );  
    }  
}
```

- Just like with the serializer, we are transforming C++ code into additional C++ code for snooper
- Something is becoming clearer now. Generating code as a bunch of printf() statements would really be awkward. Because we now have to print out serialize, unserialize, snoop and possibly collect statistics routines.

Snooper: The Same Json




```
#include "..."  
class CRnCapsuleShape :  
  public CRnShape  
{  
  public:  
  ...  
  private:  
    Vector m_vCenter[ 2 ];  
    float m_flRadius;  
    AUTO_SERIALIZE;  
};
```

```
"CRnCapsuleShape" : {  
  "fields" : {  
    "m_flRadius" : {  
      "typeName" : "float",  
    },  
    "m_vCenter" : {  
      "className" : "vector",  
      "arraySize" : 2,  
    }  
  },  
  "bases" : [  
    "CRnShape"  
  ]  
},
```

So, we still generate the same json file as we did for serialization.

Snooper: Different StringTemplate



```
void <name>::Snoop( CRnSnooper*pIn,
                  const <name> *pLocalCopy )
{
  <class.bases : {b |
                  <b>::Snoop( pIn, pLocalCopy );}>
  <class.fields.keys, class.fields.values:
    {k,v |
      <snoop_field(name=k,props=v)>
    }
  >
  <if(class.postInitMethod)>
    AfterRestore( pIn );
  <endif>
}
```

- The StringTemplate for the snooper looks like this. It's a different string template, but it's small. All my snoop code is generated from a 200-line template. These are a few of those lines. They auto-generate the code to snoop most classes in our physics engine.
- You feed it the same json file as for the serializer. It spits out snooper C++ file. It's that simple.

Snooper: Generated Code

C++

Clang

json

StringTemplate



C++

```
void <name>::Snoop( CRnSnooper*pIn,
                  const <name> *pLocalCopy )
{
    <class.bases : {b |
                  <b>::Snoop( pIn, pLocalCopy );}>
    <class.fields.keys, class.fields.values:
      {k,v |
        <snoop_field(name=k,props=v)>
      }
    >
    <if(class.postInitMethod)>
      AfterRestore( pIn );
    <endif>
}
```

```
void CRnCapsuleShape::Snoop( CRnSnooper*pIn,
                             const CRnCapsuleShape *pLocalCopy )
{
    CRnShape::Snoop( pIn, pLocalCopy );
    m_flRadius = pLocalCopy->m_flRadius;
    for( int nElement = 0;
         nElement < 2; ++nElement )
    {
        ::Snoop( pIn,
                 &pLocalCopy->m_vCenter[nElement],
                 m_vCenter[nElement] );
    }
    AfterRestore( pIn );
}
```

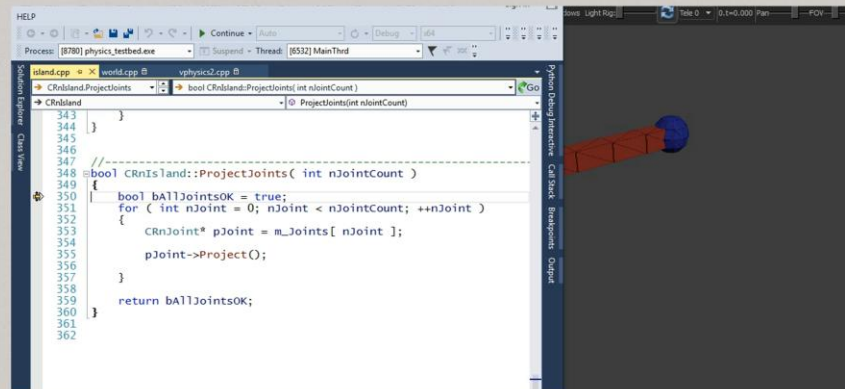
- This here is the template and the resulting generated C++ code for snooper.
- I'm snooping the whole struct into a local copy. Then I read out all the plain data elements.
- All the non-plain data elements are recursively snooped by other overloaded Snoop functions.
- The template is smart enough to know where I have arrays, templates, pointers to follow and so on.

Full Pipeline



- Just to bring the point home, here's how the full code parsing and generation pipeline looks in our case
- Parsing C++ is probably the most intimidating thing in this case. But after you spend a day or two learning Clang's API, you might find it rather intuitive.

Example: Snooping Joint Stack



The screenshot shows a debugger window with a C++ code editor on the left and a 3D visualization on the right. The code editor displays the following code:

```
343 }  
344 }  
345 }  
346 }  
347 //-----  
348 bool CRnIsland::ProjectJoints( int nJointCount )  
349 {  
350     bool bAllJointsOK = true;  
351     for ( int nJoint = 0; nJoint < nJointCount; ++nJoint )  
352     {  
353         CRnJoint* pJoint = m_Joints[ nJoint ];  
354         pJoint->Project();  
355     }  
356 }  
357 }  
358 }  
359 }  
360 }  
361 }  
362 }
```

The 3D visualization on the right shows a chain of joints, with a red and blue joint highlighted. The debugger interface includes a process window showing 'physics_testbed.exe' and a thread window showing '6532 MainThrd'.

Let's see a little demo of what we have as a result.

I configured a projection-type solver of a chain of hinge joints. As you can see, small chains solve fine, but larger chains start to break constraints.

I wanted to see how the solver behaves inside the iteration loop.

So I put a breakpoint where each joint is solved.

After each breakpoint hits, I switch to the Testbed and snoop. I have two testbeds here, one debugging the other. Notice that one is unresponsive – that's the one that I'm debugging. And another one acts like a watch window.

So what we have here effectively a visualization of the internal loop of a projection constraint solver.

Details



Fixing stuff, whether it's yours or not, is unavoidable. At least if you make anything complex.


The faster you can drill down to details, the more efficient you are while debugging problems, the more time you spend developing and not fixing stuff. Once you make debugging 10x faster and more convenient, there's a quantum leap : you can develop features with complexity that would otherwise be infeasible to develop.

That's why it's important to make the process as seamless and convenient as possible. It pays with higher productivity in the end.


There are a few more technical details, though

Divide and Conquer

Physics World



Physics Object



Physics Properties

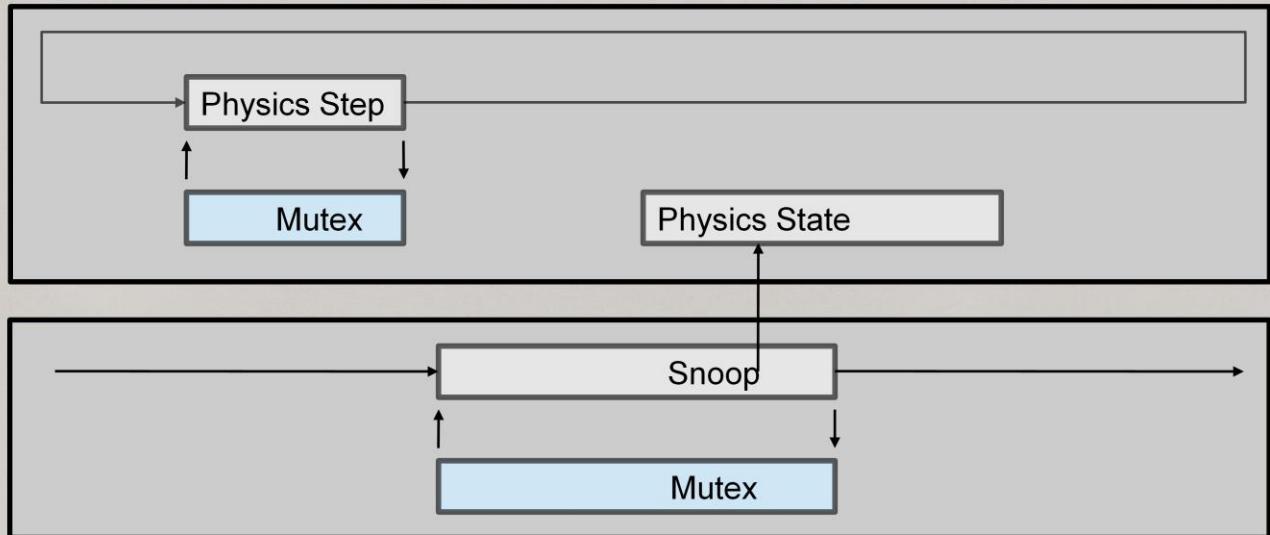


Our Testbed is a tool for divide-and-conquer debugging. It makes it quick and painless.

For example, we have a collision filtering system, and time to time something is not colliding right in the game. We just fly around, click on objects that don't collide and drill down to their flags. In most cases we see exactly what's happening within seconds, and you don't have to be a programmer to do that.

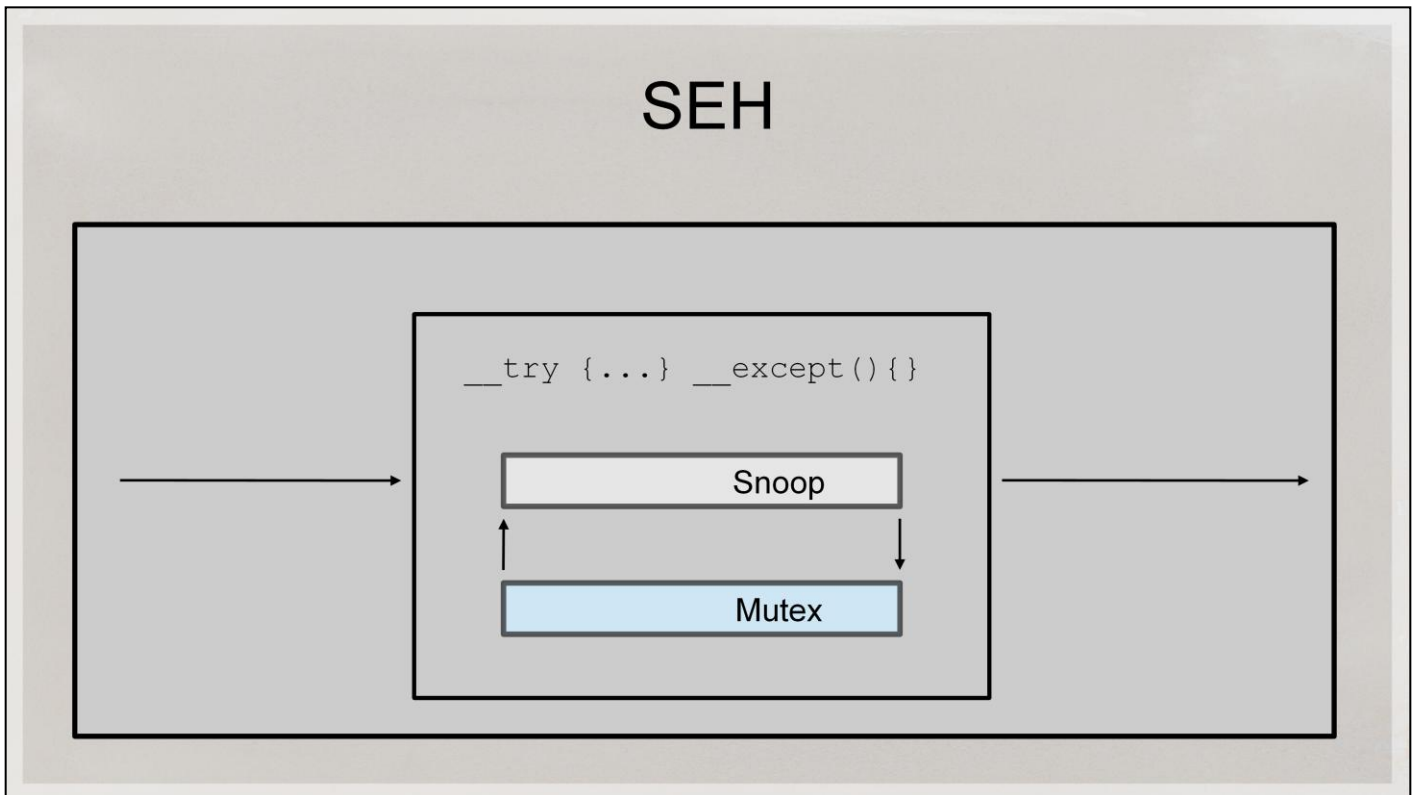
Another example: we had an incredibly slow physics. It worked fine, but was very slow. All contacts looked correctly. So we snooped and clicked the objects that were awake. Their contacts looked fine. It took us some seconds to realize there were too many contacts on them.

Mutex in game loop



- One consideration here is thread safety
- While you cannot use the Watch window in VS IDE while the game process is running, you can snoop.
- Use Mutex to guard the physics step routine

SEH



- Snooping a running process is tricky: data structures are in flux while you are reading them.
- It's impossible to crash the game by reading data from it. But the snooper might crash if you snoop an inconsistent memory. I'm just using (`__try...__except`) to catch that. Strictly speaking it's not a safe solution, but it makes the experience smoother.

SEH != try...catch

```
__try {  
    ...  
} __except(){  
    ...  
}
```

!=

```
try {  
    ...  
} catch(...){  
    ...  
}
```

*** Also, neither is necessary

Just to be clear

Structured exception handling is not the same as C++ exception handling.

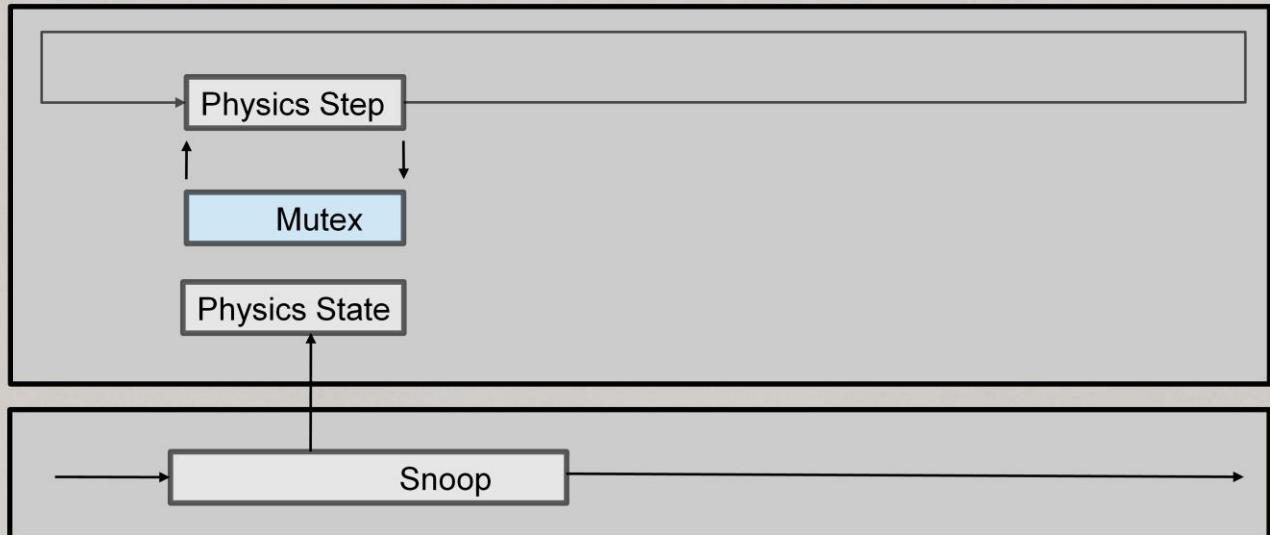
This goes outside the topic of this talk.

In any case, you don't really have to use either of these.

I just find it nice to wrap my sloppy debugging tool code in case I happen to snoop bad data structure.

If I released it as a commercial product, I'd expend the extra effort to detect that memory is unreadable and not crash.

Stopped in IDE: No Mutex



- When the game is stopped on a breakpoint, just ignore the Mutex. Or time it out.
- Data structures cannot change when the game is frozen, so there's a high chance snoop will succeed if they are consistent. It depends on where you put the breakpoint.

Clang Annotations

```
# define CLANG_ATTR(ATTR) \
    __attribute__( ( annotate( ATTR ) ) )
#define SERIALIZE_ARRAY_SIZE( SIZE ) \
    CLANG_ATTR( "array_size:" #SIZE )
...
class CSomeClass
{
...
    uint16 *m_pNodeToCtrl
        SERIALIZE_ARRAY_SIZE( m_nNodeCount );
    uint16 m_nNodeCount;
}
```

```
bool HasAnnotatedAttr(
    const clang::Decl *pDecl,
    const char * pSubstr )
{
    if( clang::AnnotateAttr *pAnnotate =
        pDecl->getAttr<clang::AnnotateAttr>() )
    {
        ShortStringRefVector attrs;
        pAnnotate->
            getAnnotation().split( attrs, " " );
        return std::find( attrs.begin(),
            attrs.end(), pSubstr ) !=
            attrs.end();
    }
    return false;
}
```

- Sometimes we need some annotations to our code. E.g. for each pointer you serializer will need to know how many objects it's pointing to. It's easy to put this in (see the left slide).
- And it's easy to add and read those annotations in the parser (see the right slide)

Complex Case: Polymorphism

```
struct VtableRecord_t
{
    template <typename T>
    void Init( CUtilStringToken name )
    {
        T* pObject = new T;
        m_nVTable = *( ( uintp* )pObject );
        delete pObject;
        m_nClassName = name.m_nHashCode;
    }

    uint64 m_nVTable;
    uint64 m_nClassName;
};
```

- I'd like to talk about a couple of less trivial cases, like classes with Vtables.
- I auto-generate a function that creates an instance of each class with Vtable, and copies its Vtable point into a known place (an array of uintptr_t).
- The snooper can then read the array and use it to recognize class type by its vtable pointer. It's like RTTI that works across process boundaries.

Complex Case: Polymorphism

```
void InitVtableRecord( CUtilVector<vtableRecord_t> *pTable )
{
    <data.classes.keys, data.classes.values: {k,v | <register_class(name=k,class=v)> }>
}

...

register_class(name,class) ::= <<
    <if(class.isLeaf)>
        ( *pTable )[ pTable->AddToTail() ].Init\<<name>\> ("<name>" );
    <else>
    <if(!class.isBase)>
        // <name> is neither leaf nor base, no need to register it for auto-recognition
    <endif>
>>
```

- This is how I generate that code with StringTemplates

Complex Case: Polymorphism

```
void InitVtableRecord( CUtilVector<vtableRecord_t> *pTable )
{
    ( *pTable )[ pTable->AddToTail() ].Init<CRnWeldJoint>( "CRnWeldJoint" );
    // CRnShadowController is neither leaf nor base, no need to register it for auto-recognition
    ( *pTable )[ pTable->AddToTail() ].Init<CRnSpringJoint>( "CRnSpringJoint" );
    ...
}
```

- And this is what the generated code looks like. One line is generated for every class with vtable. It's done automatically, so you don't have to remember to add a class to the list once you add a virtual .
- You don't have to remove a class when you remove the last virtual from it either.

More Complex Cases

Compound data structures

- World Root
- Rigid Bodies
- Joints
- Contacts
- etc.

Automated CodeGen

These data structures change frequently
Easy to generate code for
Easy to mark up with clang annotations

Primitive data types and special cases

- float, int and other PODs
- Simple types without pointers
- `std::vector<>` and external containers
- Special containers

Write code by hand

Change infrequently
Sometimes hard to generate code for
Building blocks for compound data

For some data structures, it's just easier to write serialization code by hand. The code generator should know about leafy data types like int and float, and generate serialization code automatically. There's no problem with structures only consisting of those data types, those can be simply copied directly from another process address space.


It's a bit more involved with pointers. In general, C++ does not provide enough semantics for you to know if a pointer points to one object, or a vector of objects. I'm using clang annotation to add that semantics. So, the code generator looks at that annotation and computes the size of the array and serializes accordingly.

We don't use multiple heaps for our physics engine, but if you do, you'll have to provide semantics about how to allocate the snooped or unserialized data structures.

In strange cases, like an int that is a pointer, but the lower 2 bits have a special meaning, you probably want to write that routine by hand. Also when you have a union and have to run code to decide what it means.

You have a choice whether to make your code auto-generator smarter (and harder to maintain) or snoop a specific class manually. It's a judgement call.


Statistics



***All generated code is typo-free

- We currently have 46 non-trivial classes with 320 fields for which serializer and snoop code is generated.



Statistics



***All generated code is typo-free

- The parser that uses clang is about 1000 lines long. I probably spent a couple of weeks part time writing it, and I've been using it for a year.


Statistics



***All generated code is typo-free

- The generated json is 2800 lines long


Statistics



***All generated code is typo-free

- Snooper template is 200 lines long .
- Serializer template is 400 lines.

Statistics



***All generated code is typo-free

- Auto-generated snooper is 2500 lines long.
- Serializer is 5600 lines long, because it includes 3 directions: serialize to stream, unserialize from stream, count number of bytes used (very similar to Serialize, but actually counts allocated memory).
- It takes 2-8ms to snoop a typical game frame, which is normally maybe 10 mb large

Physics in Games: Integration



Every game physics engine eventually gets integrated into a game. It makes change-compile-test iteration much longer. Especially when artists start building a huge map, and new exciting bugs show up that only happen in this huge map once in a blue moon.


It makes bug hunting mind-numbingly slow. Because the game code is much, much larger than the physics code we know so well.

Physics in Games: Perspective



Game code is written by many people over many years. Physics engine is self-contained and neat, almost miniature in comparison.
At valve, the sheer amount of logic in the old game code with history is very large.

Physics in Games: Perspective




It's really easy to lose perspective when you work on physics for a long time.

The ultimate goal of physics in a game is to make the game better. It's not to make the best or fastest or true to life physics simulation. That is why in-game debugging and optimization is very important.

If the game wants to cast a lot of rays, we need to optimize that first. Games don't generally have stacks of cubes everywhere, so that's not the best benchmark for a game physics engine. It's best to put it in the game and debug and profile it in the game, if you want it to perform well in the game.

Game Data Visualization



Run Generating_Nav

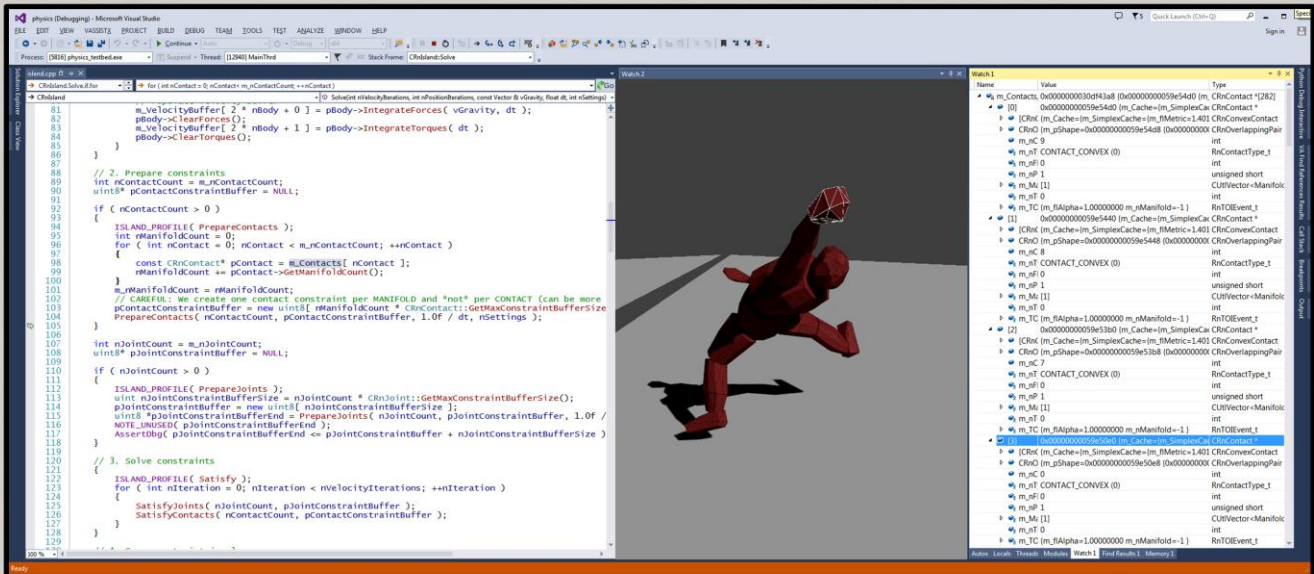
Here's another example of an unexpected benefit of visualizing everything in physics engine. Testbed records and displays all traces. And left4dead navigation system does a lot of traces. So, when connected to the game, Testbed visualizes an AI algorithm.

Set Your Priorities



I suggest we can all spend a couple of weeks up front to make our life more comfortable and deliver a better product. Sometimes you can't do that, but generally if your time horizon is years, you can afford a couple of weeks. What I'm describing is pretty general technique. Game physics is just an example, but it's usable in any complex software.

Debugging Visually



Eventually, I want my Visual Studio to look like this when I'm debugging. You can make a debugger extension with VS Extensibility, and I might make one some day. But when I experimented with it, it was very inconvenient to write a VS extension of such complexity.

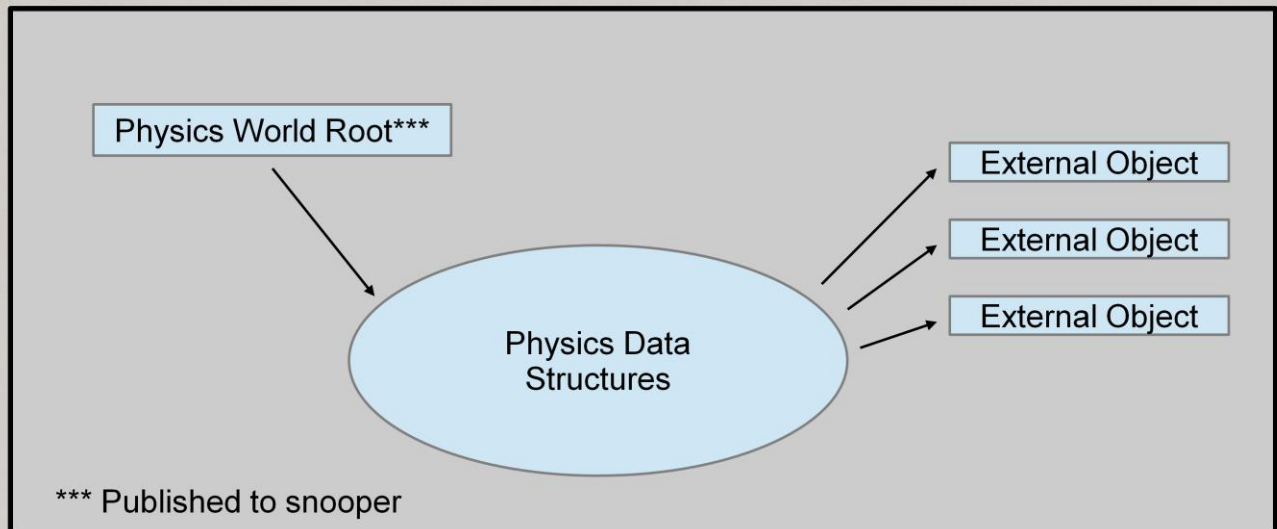
AutoExp.dat, *.natvis

```
[AutoExpand]
Vector =x=<x,g> y=<y,g> z=<z,g>
Quaternion= x=<x,g> y=<y,g> z=<z,g> w=<w,g>

[Visualizer]
CStrongHandle<*> {
    preview( $c.m_pBinding->m_Name->m_ResourceNameSymbol.u.m_pAsString )
    children (
        #(
            Data: ($T1 *)($c.m_pBinding->m_pData),
            [raw members]: [$c,!]
        )
    )
}
```

I hope everyone knows about autoexp.dat. It lets you visualize in text, so to say. I strongly suggest you just put all your data structures in there. Try to make them more readable - in your Watch window. It helps a lot.

Organizing Data Structures



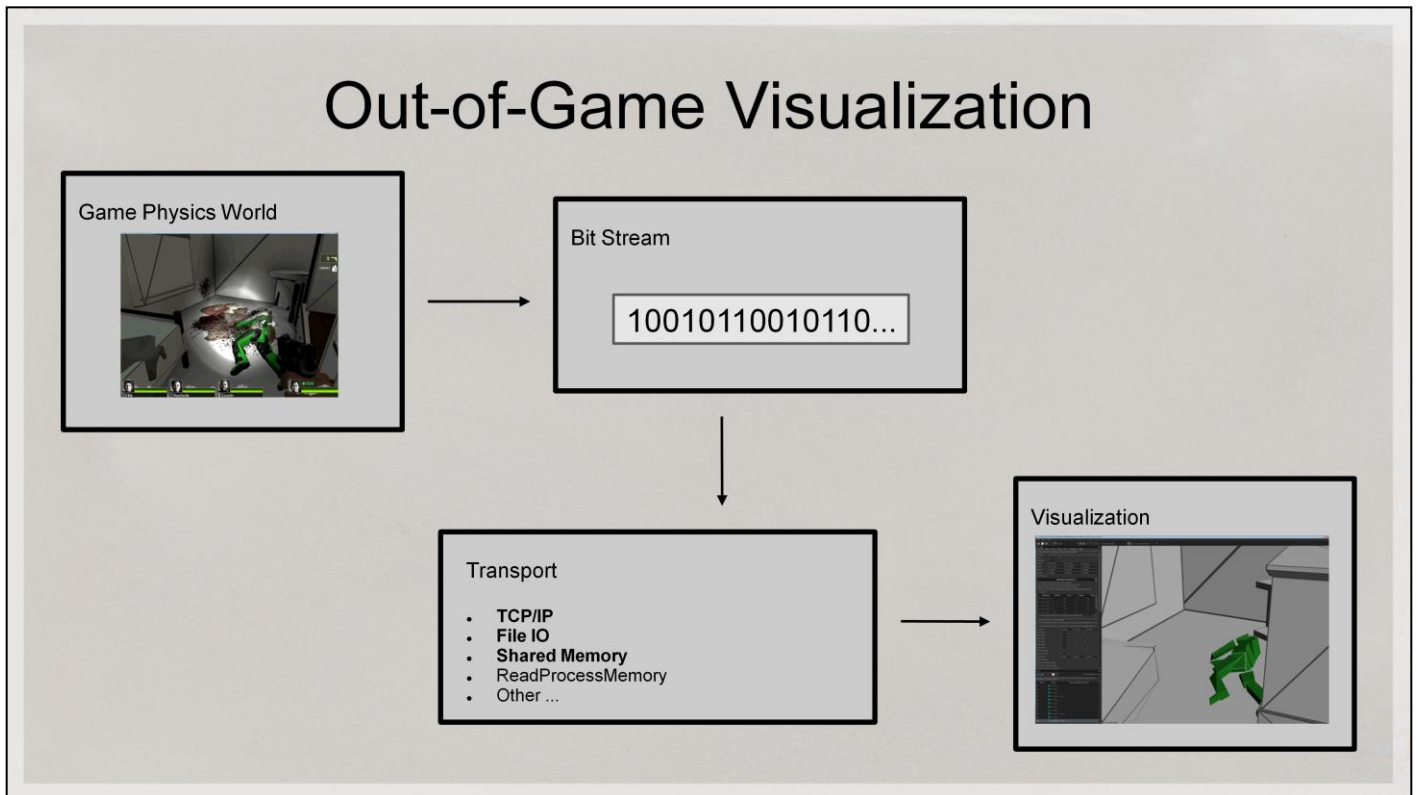
A word about organizing your data structures.

It's most convenient to have a root object from which you can crawl all your data. I just gather all the globals in one object, and only publish a pointer to that one object to the snooper.

Physics will have some data structures. You'll parse and process them. You'll probably have some external pointers, like pointers to the vertex and index buffers for debug drawing.

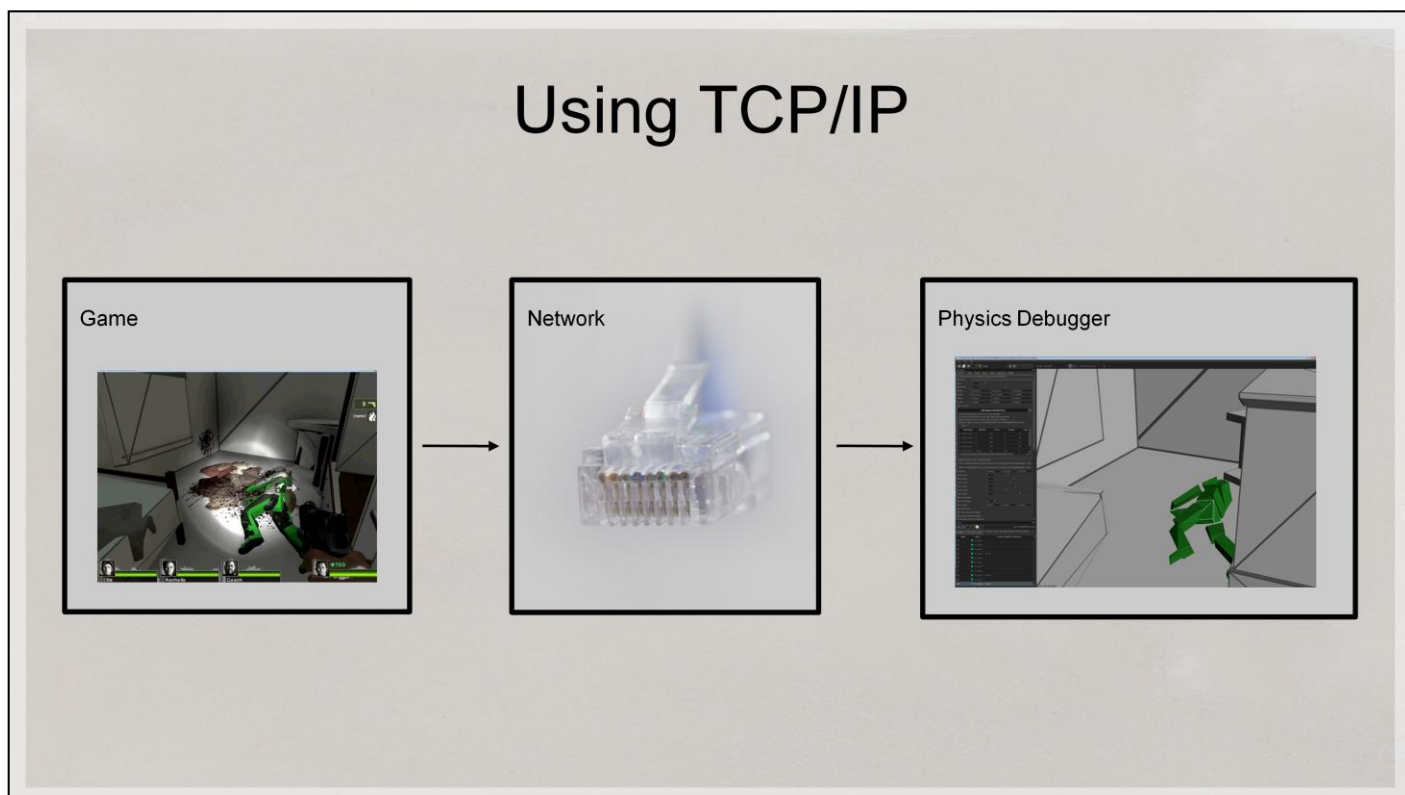
You can write routines for snooping and serialization of those things manually

Out-of-Game Visualization



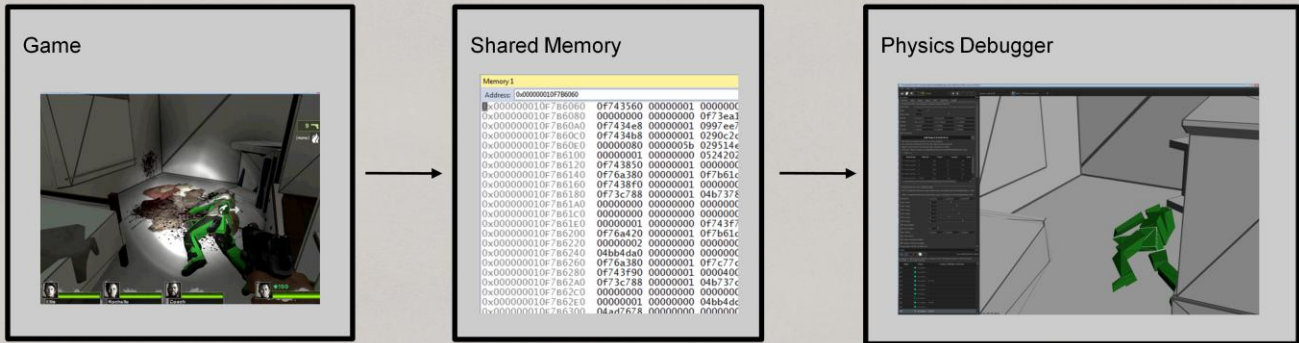
Another thought. Visualizing in an external app doesn't have to use ReadProcessMemory API.

Using TCP/IP



All the leading physics engine have some sort of over-TCP/IP visualization. Unfortunately it doesn't work when the game is stopped in debugger, frozen or crashed. Unless you stream your game state all the time, and that spends CPU. If you don't know where the problem is, and stream a lot of detailed data, it really takes a lot of CPU. And you have to remember to turn it on. But it works on remote machine. So it's a viable option.

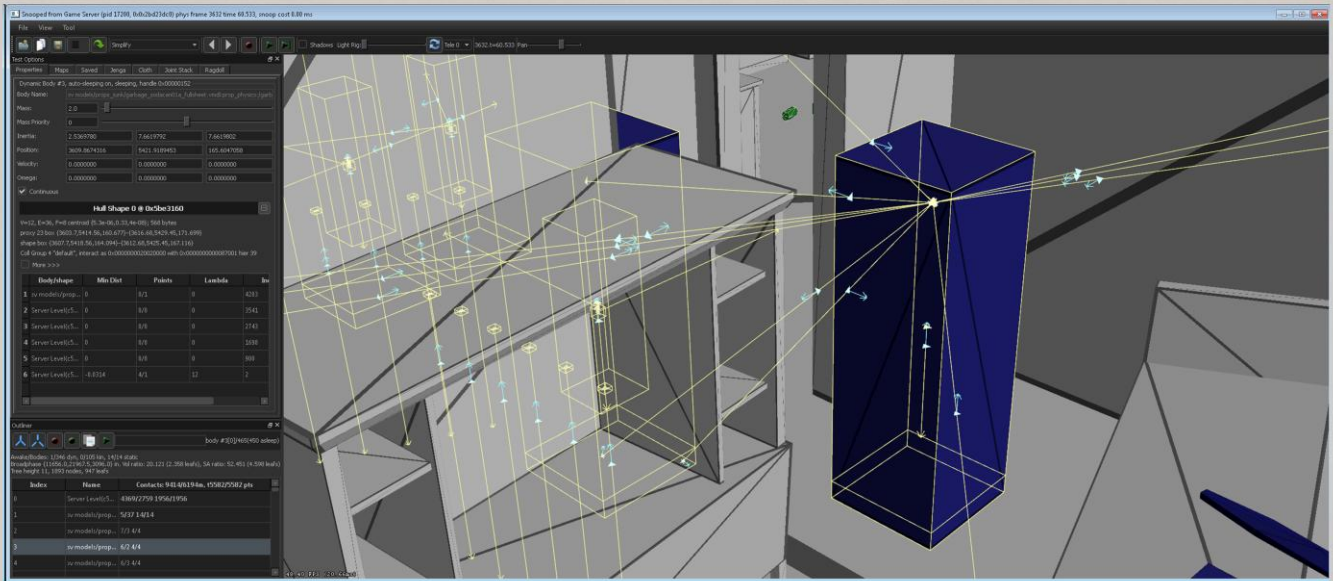
Using Shared Memory



An alternative to TCP/IP that works when the game is stopped in debugger (or frozen or crashed) would be to use shared memory.

Win32 API allows processes to share memory directly. One way is to use a memory-mapped file, but you need to dedicate a region of memory for sharing, and some of the memory would be passed down from the game (like pointer to collision resources), and it would be very inconvenient to make sure all of that is allocated in the shared memory region

Game-centric Viz



- Hopefully we'll eventually be able to save/load/rewind the whole game with this method. Our game code base is much bigger than physics engine though, so that's a lot of work.

Serialization

```
class Contact
{
    Shape* m_pShape[ 2 ];
    GraphEdge< Shape > m_Next[ 2 ];
    AUTO_SERIALIZE_BASE;
};
```

```
class Contact_Serialized
{
    int m_nShapeIndex[ 2 ];
    int m_nNextIndex[ 2 ];
};
```

10010110010110...

- We are using the same framework to generate serialize/unserialize code. Every time we change our data structures, we just re-run the parser/generator, and it's all updated
- We serialize/unserialize a byte stream. But another variant would be e.g. to generate a reflected version of all your data, with indices instead of pointers. If you have reflection API, it will let you version the serialized data

Techniques: The Summary

Technology

- Clang
- String Template
- ReadProcesMemory API

Technique

- Detailed Statistics
- Serialization
- Streaming
- Snooping
- etc.

- To summarize, we were talking about a few uses of two technologies: Clang and String Template, and one simple Windows API: ReadProcessMemory.
- Together, these technologies can enable very powerful and complete serialization, an interesting debugging technique I call snooping, gathering detailed statistics without tedious coding and without that code going stale the next time someone adds a new array somewhere.
- The same serialization template can also be used to send data over TCP/IP, and/or serialize into an alternative data structure with reflection API (e.g. into a series of dictionaries) that can be saved and loaded into newer versions of the engine, and many more uses.
- It is possible to use snooping to read and display every frame in your game
- You will probably skip some frames unless you take care to synchronize
- You will probably waste time at the mutex, for the game has to wait for the snooper to finish, and it's less efficient to snoop than to pack local game data and send it through a pipe or socket.
- It is much faster to let the game actively serialize the world delta every frame and send it over the wire. It will also guarantee no skipped frames.
- It is easier to just reuse existing serialization and send every frame over the wire
- For a quick fix, just see the biggest pieces of data (probably mesh descriptions) and cut them out of the stream if you already sent them
- When visualizing streamed physics, I tied the Vbs/IBs directly into the world data structures.
- This automatically recreated Vbs/IBs on every streamed frame
- To fix that, I had to cache those debug objects off and have special callbacks reuse them. I did it with clang's annotations.

Pros & Cons

Pros

- Reliable Process
- Repeatable Debugging
- Human Error Excluded
- Neither Tedious nor Outdated Code

Con

- Implementation Complexity


- Full-world serialization of this type includes everything, and has very high reliability. You don't have to worry about missing a thing or two: the parser/generator is unlikely to make a human error.
- This is usable to dump suspicious game frames for later examination, or streaming them. When debugging later, it's very much like debugging in-game, but it's repeatable and much more convenient.
- I don't like manual non creative work, and serialization code is always very low in creativity. Getting rid of it makes iteration faster and doesn't distract me from the creative process.
- Among the drawbacks, this method is more complex to start using it. You can still use the old methods in parallel with this method.

Conclusion



Make work fun! It pays to do that.

Q&A



Special Thanks: Content & Delivery

Erwin Coumans
Dirk Gregorius
Dennis Gustafsson
Julien Merceron

Special Thanks: Art & Style

Ricardo Ariza
Jason Brashill
Cam Fielding
Tristan Reidford

References

Clang - <http://clang.llvm.org/>

StringTemplate - <http://www.stringtemplate.org/>

ANTLR - <http://www.antlr.org/>

Valve – <http://www.valvesoftware.com/>

Steam - <http://store.steampowered.com/>

SteamDevDays - <http://www.steamdevdays.com/>

My email – sergiy@valvesoftware.com