

100% RECYCLED PAPER
PRINTED ON 100% RECYCLED PAPER



BUILDING NIGHT CITY: THE TECHNOLOGY OF CYBERPUNK 2077

8234299 - AXRC6 0001

GDC March 30-31, 2023
San Francisco, CA

Charles Tremblay | Engineering Director



BRIEF HISTORY OF RED ENGINE

REDengine 1



The Witcher 2:
Assassins of Kings

REDengine 2
Console Support



The Witcher 2:
Assassins of Kings
Enhanced Edition

REDengine 3
Open World Support



The Witcher 3:
Wild Hunt

REDengine 4



Cyberpunk 2077



March 20-24, 2022
San Francisco, CA



CYBERPUNK REQUIREMENTS



- Night City – Living metropolis in a dystopian future
 - 16x16km world and vertical
 - High density
 - No loading screen
- Vehicles
 - 140km/hr
- Scalability
 - Various platform support
 - From Xbox One to High-End PC

RED ENGINE PILLARS

- Engine / Gameplay agnostic of Editor/Tools code
 - RPC backend ↔ Editor
- Systems Scalability
 - Maximize platform hardware utilization
 - Adapt according to game state or current quest
- Quest is King



RED ENGINE RULES



No code can actively wait for anything



Code should assume it never runs in isolation



Runtime memory allocation limited to ≤ 512 bytes



No STL containers, Exception, RTTI

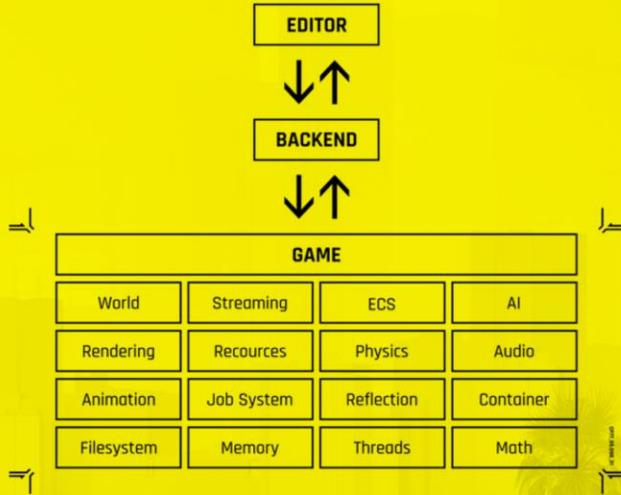


March 30-31, 2022
San Francisco, CA



CD PROJEKT RED

ENGINE - DIAGRAM



AGENDA

- Memory
- Job Systems
- Resources & IO
- Graphics
- World & Streaming
- ECS
- Systems
- Frame & Performance
- Conclusion

GDC March 30-31, 2022
San Francisco, CA



PROTOCOL 652D-A44



CYBERPUNK
2-0-7-7

MEMORY MANAGEMENT



ИГРА: ЦИБЕРПАНК 2077. АВТОР: CD PROJEKT RED. ИЗДАТЕЛЬ: CD PROJEKT RED. СЕРИЯ: ЦИБЕРПАНК 2077. СЕРИЯ: ЦИБЕРПАНК 2077. СЕРИЯ: ЦИБЕРПАНК 2077.

MEMORY – REQUIREMENTS

- Allocators for specific needs
- Default allocator is lockless while minimizing fragmentation
- Easy to use and understand
- Easy to extend
- Budgets are easy to define
- Every allocator needs to be under the proper budgets
- Reliable



March 30-31, 2022
San Francisco, CA



MEMORY – ALLOCATOR FOR SPECIFIC NEEDS

- Slab
- TLF
- Fixed Size
- Linear
- Buddy
- Stack
- Job & Frame Allocator
- And more ... !



March 30-31, 2022
San Francisco, CA



MEMORY – DEFAULT ALLOCATOR

- Custom Slab allocator with explicit thread registration for $\leq 512b$ allocations
 - From Witcher 3 experiences: 75% + of all allocations are less than 512b
 - Even distribution across our job/task threads
 - Good locality
 - 10ns on PS4
 - Average waste of 6% per thread
- TLSF allocator for allocations between 512b a 512kb
- “BigSize” allocator for allocations $> 512kb$

MEMORY – EASY TO EXTEND

- Allocator code does not have to reside in memory project
- Clear and minimal static interface to fulfill
- All utilities provided by memory system are available

```

struct SimpleAllocatorMetrics();

class SimpleAllocator
{
public:
    RED_MEMORY_DECLARE_ALLOCATOR( SimpleAllocator, SimpleAllocatorMetrics, 16 );

    red::memory::Block Allocate( uint32_t size );
    red::memory::Block AllocateAligned( uint32_t size, uint32_t alignment );
    red::memory::Block Reallocate( red::memory::Block& block, uint32_t size );
    red::memory::Block ReallocateAligned( red::memory::Block& block, uint32_t size, uint32_t alignment );
    void Free( red::memory::Block& block );
    uint64_t GetBlockSize( uint64_t block ) const;
    void SerializeMetrics( red::memory::Serializer& serializer );
};
    
```

MEMORY – EASY TO USE & UNDERSTAND

- Fully documented
- Code consistency
- Easy to read!

```
////////////////////////////////////  
// 3.2 Allocating from a specific Allocator  
// Like with the new operator replacement (RED_NEW), you can also provide an allocator explicitly.  
// For example:  
void Sample_3_2_Allocator()  
{  
    const int32_t allocatorBufferSize = RED_KILO_BYTE( 64 );  
    void * allocatorBuffer = RED_ALLOCATE( red::PoolDefault, allocatorBufferSize );  
    const red::memory::StaticTLSAllocatorParameter param = { allocatorBuffer, allocatorBufferSize };  
    red::memory::StaticTLSAllocator allocator;  
    allocator.Initialize( param );  
  
    void * buffer = RED_ALLOCATE( allocator, 128 );  
    void * reallocBuffer = RED_REALLOCATE( allocator, buffer, 256 );  
    RED_FREE( allocator, buffer );  
    RED_FREE( allocator, reallocBuffer );  
}
```

MEMORY – POOLS

- All allocations needs to be associated to a Pool
- Pools define budgets
- Pools can be parented

```
RED_MEMORY_POOL( PoolAI_Behaviour, red::memory::DefaultAllocator, AI_API );

void MemoryPoolSnippet()
{
    auto& allocator = red::memory::AcquireDefaultAllocator();
    RED_INITIALIZE_MEMORY_POOL( PoolAI_Behaviour, AI::PoolAI, allocator, RED_MEGA_BYTE( 7 ) );

    auto * scalar = RED_NEW( int32_t, PoolAI_Behaviour )( 123 );
    RED_DELETE( scalar, PoolAI_Behaviour );

    void* buffer = RED_ALLOCATE( PoolAI_Behaviour, RED_KILO_BYTE( 64 ) );
    RED_FREE( PoolAI_Behaviour, buffer );

    red::DynArray< int32_t > myArray( PoolAI_Behaviour() );
    myArray.Reserve( 16 );
}
```

MEMORY – METRICS & TRACKING

- All memory allocations can be tracked
- Report can be used for automated tools

name	usage	include	include count	exclusive	exclusive count	peak	budget	inc. alloc. bytes/frame	inc. alloc. count/frame	debug	alloc	contributes to parent	handles	parenting
Parent Pools														
Selected Pool:														
PoolCPU	100%	2.34 GB	842777	0 B	0	0 B	1.54 GB	4.60 MB	1104	no	yes		0x11305107	yes
Children Pools														
PoolCriticalSection	54%	1.68 GB	9	1.64 GB	3	1.64 GB	4.07 KB	0 B	0	no	yes		0x71a63009	yes
Pool	3%	122.11 MB	184283	816.89 KB	1844	817.00 KB	140.01 MB	128 B	8	no	yes		0x01504114	yes
PoolRenderTarget	3%	171.16 MB	129239	10.04 MB	19813	10.04 MB	195.01 MB	2.24 KB	41	no	yes		0x07880263	yes
PoolTexture	31%	196.65 MB	492610	22.71 MB	1310	26.70 MB	195.01 MB	65.51 KB	81	no	yes		0x07880263	yes
PoolCompression	0%	0 B	0	0 B	0	0 B	1.02 KB	0 B	0	no	yes		0x18282428	yes
PoolVertexBuffer	14%	1.48 MB	8950	1.48 MB	8950	1.56 MB	1.00 MB	0 B	0	no	yes		0x10424174	yes
PoolVertexBufferFrame	0%	7.40 KB	1	7.40 KB	1	82.06 KB	32.00 MB	32.25 KB	22	no	yes		0x71a63009	yes
PoolEngine	18%	791.08 MB	444180	64.18 MB	24437	81.01 MB	432.01 MB	101.41 KB	8729	no	yes		0x01504114	yes
PoolFrame	0%	0 B	0	0 B	0	210.29 KB	10.00 MB	64.29 KB	102	no	yes		0x01504114	yes
PoolRenderTargetData	0%	420 B	5	420 B	5	420 B	1.00 MB	0 B	0	no	yes		0x11305107	yes
PoolTexture	17%	249.31 MB	161766	0 B	0	0 B	145.01 MB	4.80 KB	91	no	yes		0x01504114	yes
PoolTexture	0%	92.13 KB	28	24 B	1	360 B	1.00 MB	0 B	0	no	yes		0x01504114	yes
PoolImage	0%	0 B	0	0 B	0	405.53 KB	1.00 KB	0 B	0	no	yes		0x01504114	yes
PoolLayerGenerator	1%	391.41 KB	7133	391.41 KB	7133	199.17 KB	1.00 MB	240 B	10	no	yes		0x01504114	yes
PoolMultiLayer	0%	984 B	3	0 B	0	0 B	220.01 MB	0 B	0	no	yes		0x01504114	yes
PoolMultiLayerCanvas	0%	0 B	0	0 B	0	0 B	64.00 MB	0 B	0	no	yes		0x01504114	yes
PoolTextureService	0%	0 B	0	0 B	0	0 B	2.00 MB	0 B	0	no	yes		0x71a63009	yes
PoolPhysics	3%	91.72 MB	101149	4.32 MB	1291	4.37 MB	100.00 MB	99.00 KB	68	no	yes		0x01504114	yes
PoolRendering	21%	716.36 MB	779634	0 B	0	0 B	285.00 MB	3.59 MB	1011	no	yes		0x01504114	yes
PoolGames	2%	1.47 MB	25428	77.36 KB	527	77.36 KB	24.00 MB	1.04 KB	12	no	yes		0x11305107	yes
Pool	3%	103.94 MB	79012	0 B	0	0 B	15.00 MB	383.88 KB	202	no	yes		0x01504114	yes

MEMORY – BUDGETS

CPU - 1.5gb

- Rendering – 300mb
- Animation – 200mb
- Audio – 200mb
- Streaming – 160mb
- Gameplay – 160mb
- AI – 140mb
- Resources – 100mb
- Physics – 100mb
- Archives – 64mb
- UI – 55mb

GPU - 3gb

- Texture Generic – 700mb
- Texture Multilayer – 350mb
- Render Targets – 640mb
- Mesh – 700mb
- GI – 300mb



March 20-24, 2022
San Francisco, CA



CONTENT BUDGET

194mb Environment Textures	33mb Weapon Textures	33mb Vehicle Textures	35mb FX Textures	221mb Code Based Textures	72mb Surface Textures
107mb Environment Meshes	50mb Weapon Meshes	50mb Vehicle Meshes	10mb FX Meshes	151mb Code Based meshes	210mb Multilayer library
42mb Environment MLMasks	10mb Weapon MLMasks	20mb Vehicle MLMasks	5mb FX MLMasks	36mb Code Based MLMasks	50mb Microblends
217mb Character Textures	30mb UI Textures	30mb Vehicle Proxies			
57mb Character Meshes	50mb UI Advertisements				
27mb Character MLMasks	10mb Videos				



March 20-26, 2022
San Francisco, CA





There was a good question that I feel I failed to explain correctly.

How Jobs are better than custom thread.

I can answer this with a real example.

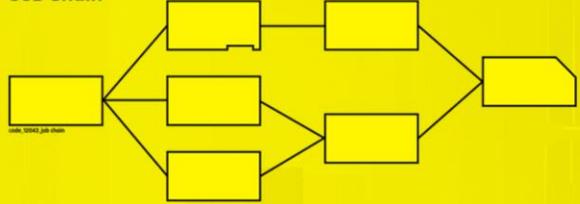
On Witcher3, the render thread will issue the culling work to be done on another thread. In meantime, render thread will continue to do its work but at one point it will need to wait for the result of the culling. With some luck it's finished, and result is available. However, very often it is not, causing the render thread to completely stop until result is finally available. In some case this cause up to 5ms completely wasted on render thread.

Now with proper job chain, this problem never occur. See Graphics part about render graph 😊

JOB SYSTEM – REQUIREMENTS

- Unshackle Main & Render thread
- Everything should be a job
No more custom threads
- Easy to build a job chain
- Easy to write continuation jobs
- Easy to use

Job Chain



JOB BUILDER

- Main utility to dispatch and manage jobs chain
- Allows to create complex job chain
- Used by every single system

```
class HeavyObject
{
public:
    void Uninitialize() {}
};

void FireAndForgetSnippet( red::UniquePtr< HeavyObject > object )
{
    job::Builder builder( job::Priority::Latent );
    builder.DispatchJob( "Uninitialize_and_Destroy_Object", [obj = std::move(object)]( const job::RunContext& )
    {
        obj->Uninitialize();
        obj.Reset();
    } );
}
```



March 30-31, 2022
San Francisco, CA



JOB BUILDER – SIMPLE JOB CHAIN

DispatchJob creates dependant job by default



```
class Object
{
public:
void FirstJob();
void SecondJob();
void ThirdJob();
};

void SimpleJobDependencySnippet( const red::SharedPtr< Object >& object )
{
job::Builder builder( job::Priority::CriticalPath );
builder.DispatchJob( "First_Job", [object]( const job::RunContext& ) {
; object->FirstJob();
} );

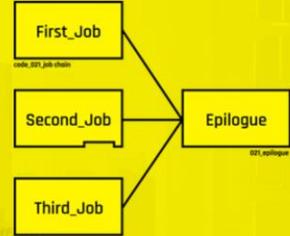
builder.DispatchJob( "Second_Job", [object]( const job::RunContext& ) {
; object->SecondJob();
} );

builder.DispatchJob( "Third_Job", [object]( const job::RunContext& ) {
; object->ThirdJob();
} );
};
```

JOB BUILDER – PARALLEL JOB CHAIN

- DispatchJob can create jobs to be run in parallel
- DispatchParallelForJob can also be used

```
void ParallelJobWithEpilogueJobSnippet( const red::SharedPtr< Object >& object )  
{  
    job::Builder builder( job::Priority::CriticalPath );  
    builder.DispatchJob< job::Fence::None >( "First_Job", [object]( const job::RunContext& ) {  
        object->FirstJob();  
    } );  
    builder.DispatchJob< job::Fence::None >( "Second_Job", [object]( const job::RunContext& ) {  
        object->SecondJob();  
    } );  
    builder.DispatchJob< job::Fence::None >( "Third_Job", [object]( const job::RunContext& ) {  
        object->ThirdJob();  
    } );  
    builder.DispatchFenceExplicitly();  
    builder.DispatchJob( "Epilogue_Job", [object]( const job::RunContext& ) {  
        object->EpilogueJob();  
    } );  
}
```





JOB SYSTEM – CANCELLING JOBS?

- It is not possible to safely cancel a complete job chain
- However it can be done on the user side

26

ИЗДАТЕЛЬСТВО «ИТ» 2017

CYBERPUNK
2-0-7-7

IO & RESOURCE MANAGEMENT



RESOURCE LOADER

- Resource loading request can be made virtually anywhere
- You cannot actively wait for completion
- With great power comes great responsibility

```
void ResourceLoadingSnippet( res::ResourceLoader* resourceLoader )  
{  
    res::IssueLoadingRequestParameter param;  
    param.path = RED_CONST_RESOURCEPATH( "base\\gameplay\\devices\\vending_machines\\vending_machine_1.ent" );  
    param.priority = io::eAsyncPriority_Background;  
  
    res::ResourceTokenHandle token = resourceLoader->issueLoadingRequest( param );  
    if( token->isLoading() )  
    {  
        // Resource was most likely already loaded. You can use right away!  
        auto& resource = token->getResource();  
    }  
    else if( token->hasFailed() ) // Resource loading request failed?  
    {  
        auto errorType = token->getError(); // It could be invalid path, extension or resource wasn't found, etc..  
    }  
    else  
    {  
        /* Resource loading request on going! */  
    }  
  
    // If refcount goes to 0, resource will be schedule for unload, or loading request will be cancelled.  
    token.Reset();  
}
```



March 20-24, 2022
San Francisco, CA



CONTINUATION JOB WHEN RESOURCE LOADED

- It is possible to link job to a resource-loading job chain
- But it's important to validate result of request!

```
void ResourceLoadingAndContinuationJobSnippet( res::ResourceLoader* resourceLoader )  
{  
    res::IssueLoadingRequestParameter param;  
    param.path = RED_CONST_RESOURCEPATH( "base\\gameplay\\devices\\vending_machines\\vending_machine_1.ent" );  
    param.priority = 10::eAsyncPriority_Background;  
  
    res::ResourceTokenHandle token = resourceLoader->IssueLoadingRequest( param );  
    job::Builder builder( job::Priority::Latent );  
    builder.DispatchWait( token->GetWaitCounter() );  
    builder.DispatchJob( "OnResourceLoadingRequestCompleted", [token]( const job::RunContext& ) {  
  
        if( !token->HasFailed() )  
        {  
            // Resource is ready! It can be use safely.  
            auto resource = token->GetResource();  
            // Note! if resource ownership is not taken, token could trigger unload when refcount will go to 0.  
        }  
    } );  
}
```



RESOURCE REQUEST UNDER HEAVY CONTENTION

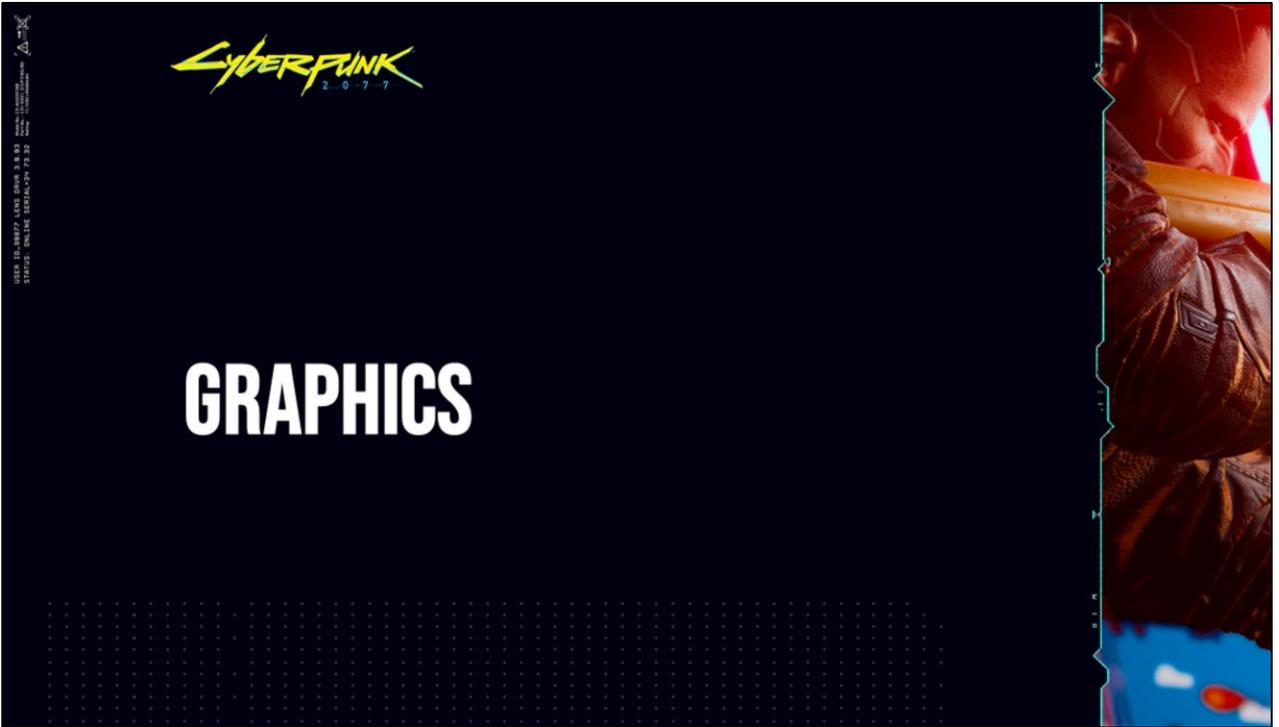
- Multiple concurrent requests
- Resources could be requested to be unloaded at the same time
- Avoid locks as much as possible



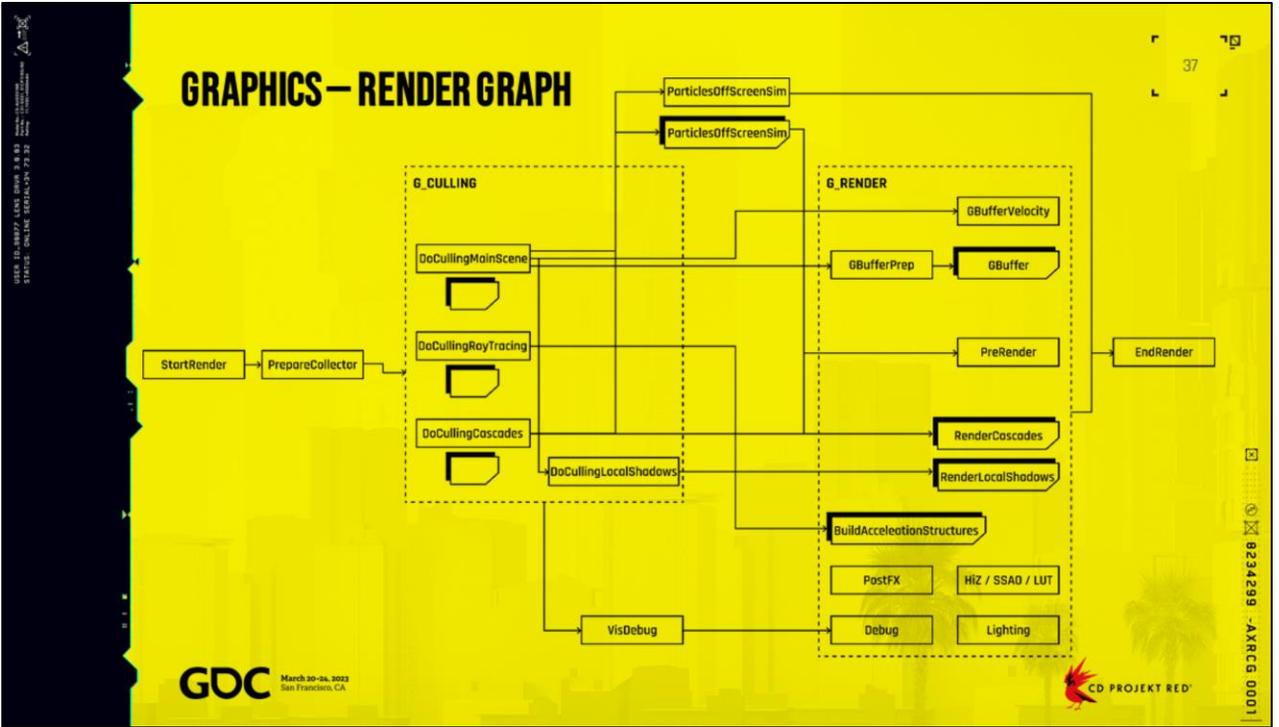
8234299 - AXRC6 0001

RESOURCE REQUEST UNDER HEAVY CONTENTION

```
ResourceTokenHandle ResourceLoader::TrySchedulingLoadingJobs( const res::ResourcePath resolvedPath, const IssueLoadingRequestParameter& param )  
{  
    ResourceTokenHandle token;  
    Job::CompletionDeferral deferral;  
    bool isNewToken = false;  
  
    RED_SCOPE_SHARED_LOCK( m_resourceTokenLock ); // SHARED LOCK: Try to acquire already created token  
    token = TryAcquireResourceToken_NoLock( resolvedPath );  
  
    if( !token ) {  
        // Create new resource token outside the lock.  
        auto tokenDeferralPair = CreateResourceToken( resolvedPath );  
  
        RED_SCOPE_LOCK( m_resourceTokenLock ); // EXCLUSIVE LOCK: Try to create new token  
        token = TryAcquireResourceToken_NoLock( resolvedPath ); // Got lock. Did someone beat us to it ?  
        if( !token ) { // If we got here, we have the lock, no one managed to beat us to it also. Insert safely  
            token = std::move( tokenDeferralPair.first );  
            token->Internal_SetPriority( param.priority );  
            deferral = std::move( tokenDeferralPair.second );  
            isNewToken = true;  
            m_resourceTokenDictionary[ resolvedPath ] = token;  
        }  
    }  
  
    if( !isNewToken ) {  
        // If we got here, we have token and we created it. First, is the Resource already loaded?  
        const Thandle< CResource > resource = m_resourceBank->FindResource( resolvedPath );  
        if( !resource ) {  
            ScheduleLoadingJob( token, std::move( deferral ), resolvedPath, param ); // No? Kickstart loading job!  
        }  
        else {  
            token->Internal_AssignLoadedResource( resource, std::move( deferral ) ); // Yes? Assign to token.  
        }  
    }  
    return token;  
}
```

Very good presentation by Tim Green, SIGGRAPH 2021



See Samples section at the end. I've added full code example on how it looks like in code

GRAPHICS - GPU FRAME



© 2017 CD Projekt S.A. All rights reserved. CD Projekt S.A. is a registered trademark of CD Projekt S.A. Cyberpunk 2077 is a trademark of CD Projekt S.A.

CYBERPUNK
2077

WORLD & STREAMING





CYBERPUNK 2077 IN NUMBERS

- One world, 16x16 km
- 15 million + Objects / Nodes / Entities
- 30 million + Foliage instances
- 31-38 main quests, 80 side quests, 74 gig quests
- 100+ NCPD Scanner Hustles, hidden gems, mini stories, & other small content pieces
- 2200+ Quest prefabs



HOW WERE PRECISION ISSUES RESOLVED?

- Fixed Point is your friend!
 - Int32, 15 / 17 for our World Position
- Physics scene origin needs to be updated
 - Every 1024m from last origin update
 - Every physics proxy needs to be updated
- Camera translation as an origin for rendered objects



WORLD NODE & NODEINSTANCE

- Nodeinstances are the units that are streamed in
- Nodeinstances are NOT updated directly
- Node is the payload provided to an instance when streaming in
- Node can be shared to multiple node instances

```
struct SimpleMeshNodeInstance : public NodeInstance
{
    RTTI_DECLARE_TYPE( SimpleMeshNodeInstance );
    virtual bool OnInitialize( const Context& context ) override final;
    virtual void OnAttach( RuntimeScene& scene ) override final;
    virtual void OnDetach( RuntimeScene& scene ) override final;

    RenderProxyPtr m_renderProxy;
    MeshResourceHandle m_loadedMesh;
    ResourceTokenHandle m_loadingToken;
};

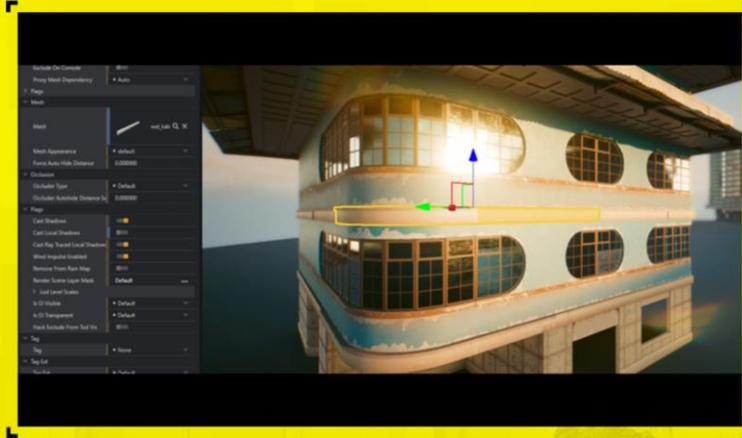
struct SimpleMeshNode : public Node
{
    RTTI_DECLARE_TYPE( SimpleMeshNode );
    virtual const rtti::ClassType* GetInstanceClass() const override final
    {
        return ClassID< SimpleMeshNodeInstance >();
    }
    TResAsyncRef< Mesh > mesh;
};
```



March 20-26, 2022
San Francisco, CA



PREFABS EDITOR



GDC

March 30-31, 2022
San Francisco, CA



STREAMING GRID— OPTIMIZED

- Sector dimensions are much smaller
 - 64m cube for exteriors
 - 32m cube for interiors
- Sectors are rebalanced to eliminate “almost empty” sectors
 - Nodes will always “move up” to higher level sector
- Quest sectors are now merged into a single sector per quest
- Resources are now embedded in sectors
 - Minimap
 - Simplified for distance mesh (we called them proxy)
 - Foliage
- Instancing nodes are generated replacing Mesh nodes using same mesh
- Runtime cost of node payload on optimized grid was around 60-80mb

RUNTIME STREAMING PROCESS

1. **Compute which sector needs to be loaded / unloaded**
 - a. Request Async Load on sector in range
 - b. Remove sector from grid that needs to be unloaded
 - c. Add loaded sector to grid
2. **Compute which nodes that need to be streamed-in / streamed-out**
 - a. Request each node in range to start streaming
 - b. Cancel streaming for nodes that aren't in range anymore
 - c. Accumulate nodes that are ready to be attached / detached
3. **Attach / detach nodes**

COMPUTE SECTORS IN RANGE

48

```
__m128 xxxx = _mm_load_ps( @position.X ), yyyy = _mm_load_ps( @position.Y ), zzzz = _mm_load_ps( @position.Z );
UInt32 index = 0;
while( index + 8 < sectorCount )
{
    char firstMask = 0;
    char secondMask = 0;
    {
        __m128 vectorMinX = _mm_load_ps( streamMinX + index ); __m128 compareMinX = _mm_cmpgt_ps( xxxx, vectorMinX );
        __m128 vectorMinY = _mm_load_ps( streamMinY + index ); __m128 compareMinY = _mm_cmpgt_ps( yyyy, vectorMinY );
        __m128 vectorMinZ = _mm_load_ps( streamMinZ + index ); __m128 compareMinZ = _mm_cmpgt_ps( zzzz, vectorMinZ );

        __m128 vectorMaxX = _mm_load_ps( streamMaxX + index ); __m128 compareMaxX = _mm_cmpgt_ps( xxxx, vectorMaxX );
        __m128 vectorMaxY = _mm_load_ps( streamMaxY + index ); __m128 compareMaxY = _mm_cmpgt_ps( yyyy, vectorMaxY );
        __m128 vectorMaxZ = _mm_load_ps( streamMaxZ + index ); __m128 compareMaxZ = _mm_cmpgt_ps( zzzz, vectorMaxZ );

        __m128 resultX = _mm_or_ps( compareMinX, compareMaxX );
        __m128 resultY = _mm_or_ps( compareMinY, compareMaxY );
        __m128 resultZ = _mm_or_ps( compareMinZ, compareMaxZ );
        __m128 results = _mm_or_ps( resultX, resultY );
        firstMask = ~_mm_movemask_ps( _mm_or_ps( results, resultZ ) );
    }

    index += 4;

    // repeat same code than above for the secondMask. Omitted for space reason.

    // Combine result into one byte, assign to bitset.
    char combinedMask = secondMask << 4;
    combinedMask |= (firstMask & 0xf);
    *sectorMaskStream |= combinedMask;
    ++sectorMaskStream;
    index += 4;
}
```

GDC

March 20-26, 2022
San Francisco, CA

CD PROJEKT RED

8234299 -AXRC6 0001

Orbis number
700us simple point contains in box
300us AOS
150us SOA

COMPUTE NODES IN RANGE

49

```
const UInt32 c_proxyStep = 16 * 1024;
for( UInt32 firstProxyIndex = 0, end = proxyCount; firstProxyIndex < end; firstProxyIndex += c_proxyStep )
{
    const UInt32 proxyCount = std::min( c_proxyStep, end - firstProxyIndex );
    builder.DispatchJob( job::Fence::None >( "Streaming_CollectNodeParallel", [0]( const job::RunContext& context ) {

        const float* streamX = m_x + firstProxyIndex, streamY = m_y + firstProxyIndex, streamZ = m_z + firstProxyIndex;
        const float* streamRadius2 = m_r2 + firstProxyIndex;
        const __m128 streamingDistanceScale128 = _mm_set1_ps( distanceMultiplier );
        __m128 xxxx = _mm_load_ps( &position.X ), yyyy = _mm_load_ps( &position.Y ), zzzz = _mm_load_ps( &position.Z );
        UInt8* maskStream = outMask + ( firstProxyIndex / 8 );
        UInt32 index = 0;

        while( index + 8 <= proxyCount )
        {
            char firstMask = 0, secondMask = 0;

            __m128 vectorX = _mm_sub_ps( _mm_load_ps( streamX + index ), xxxx ); __m128 vectorX2 = _mm_mul_ps( vectorX, vectorX );
            __m128 vectorY = _mm_sub_ps( _mm_load_ps( streamY + index ), yyyy ); __m128 vectorY2 = _mm_mul_ps( vectorY, vectorY );
            __m128 vectorZ = _mm_sub_ps( _mm_load_ps( streamZ + index ), zzzz ); __m128 vectorZ2 = _mm_mul_ps( vectorZ, vectorZ );
            __m128 result = _mm_add_ps( _mm_add_ps( vectorX2, vectorY2 ), vectorZ2 );
            __m128 vectorRadius2 = _mm_mul_ps( streamingDistanceScale128, _mm_load_ps( streamRadius2 + index ) );
            firstMask = _mm_movemask_ps( _mm_sub_ps( result, vectorRadius2 ) );

            index += 4;
            // repeat same code than above for the secondMask. Omitted for space reason.

            char combinedMask = secondMask << 4;
            combinedMask |= ( firstMask & 0xf );
            *maskStream = combinedMask;
            ++maskStream;
            index += 4;
        }
    } );
};
```

More or less 50us per job on PS4

CYBERPUNK
2-0-7-7

ECS



ИГРА: ЦИБЕРПАНК 2077. © 2013 CD PROJEKT S.A. WYDAWCA: CD PROJEKT S.A. WARSZAWA, POLSKA. WYDAWCA: CD PROJEKT S.A. WARSZAWA, POLSKA. WYDAWCA: CD PROJEKT S.A. WARSZAWA, POLSKA.

ECS — RULES

- Reserved for more complex composition
- Can be spawned at runtime
- Spawn cannot be sync
- Cannot actively wait on spawn completion
- Visuals should be decoupled from logic if possible
- Entity / component update logic should be managed by proper systems
- Entity / component cannot communicate directly to other instance

ECS—EVENT



```
struct PhysicalImpulseEvent : public red::Event
{
    RTTI_DECLARE_TYPE( PhysicalImpulseEvent );
    uint32_t m_bodyIndex;
    Vector3 m_worldImpulse;
    Vector3 m_worldPosition;
    float m_radius;
    uint32_t m_shapeIndex;
    physics::ProxyID m_proxyId;
};

RTTI_BEGIN_TYPE( MeatBag );
RTTI_PARENT_TYPE( game::Object );
RTTI_PROPERTY_CATEGORY( "Physics" );
RTTI_PROPERTY( m_kinematicBodyBoneName ).editable();
RTTI_PROPERTY( m_bagBodyBoneName ).editable();
RTTI_PROPERTY( m_physicalComponentName ).editable();
RTTI_PROPERTY( m_bagHitComponentName ).editable();
RTTI_PROPERTY( m_bagDestroyComponentName ).editable();
RTTI_PROPERTY_CATEGORY( "Effects" );
RTTI_PROPERTY( m_destructionEffectName ).editable();
RTTI_PROPERTY( m_jiggleEffectName ).editable();
RED_EVENT_CONNECTOR( OnSetup );
RED_EVENT_CONNECTOR( OnControl );
RED_EVENT_CONNECTOR( OnPhysicalImpulse );
RTTI_END_TYPE();
```

ECS—EVENT

```
void MeatBag::OnPhysicalImpulse( const PhysicalImpulseEvent& evt )
{
    if( evt.m_bodyIndex == m_bagBodyIndex ) {
        ent::SpawnEffectSetup effectSetup;
        if( --m_hitPoints > 0 ) {
            effectSetup.effectName = m_jiggleEffectName;
        }
        else {
            m_physicalComponent->ToggleCollisions( false, m_bagBodyIndex );
            m_physicalComponent->ToggleQueries( false, m_bagBodyIndex );
            m_bagHitComponent->Enable( false );
            m_bagDestroyComponent->SetForceInitAsVisible( true );
            m_bagDestroyComponent->Enable( true );
            m_bagDestroyComponent->SetForceInitAsVisible( false );
            effectSetup.effectName = m_destructionEffectName;
        }

        QueueEvent( CreateHandle< ent::SpawnEffectEvent >( effectSetup ) )
    }
}
```

ECS – ASYNC SPAWNING

- There is no sync spawn request
- However it does not need to be attached to world
- Scheduling to attach to world is guaranteed to be done before next frame start



March 30-31, 2022
San Francisco, CA



ECS — ASYNC SPAWNING

```
const RuntimeScene* scene = AcquireRuntimeScene();
RuntimeSystemEntity* entitySystem = scene->GetSystem< RuntimeSystemEntity >();
ent::EntitySpawnService* spawnService = entitySystem->AcquireSpawnService();

ent::EntityStaticSpawnContext context = {
    m_entityTemplate.GetPath(),
    m_transform,
    m_globalId,
    m_appearanceName,
    m_instanceData,
    editorService,
    ent::EntityLODInitialSetup( m_entityLod ),
    m_ioPriority
};

ent::EntitySpawnTokenHandle token = spawnService->SpawnStaticEntity( context );
builder.DispatchWait( token->GetWaitCounter() );

builder.DispatchJob( "OnSpawnReady", [token, entitySystem]( const job::RunContext& ) {
    if( !token->IsCancelled() && !token->HasFailed() )
    {
        auto entity = token->ExtractSpawnedEntity();
        entitySystem->ScheduleEntityAttach( { entity } );
    }
});
```


CYBERPUNK
2-0-7-7

SYSTEMS



CD PROJEKT S.A. CYBERPUNK 2077

SYSTEMS

60

- World Systems – lifetime limited to world
 - Also available for editor preview
- Game System – available during the whole game process
- Main method to register to frame update
- Systems can communicate with each other
 - Public interface thread safety needs to be considered

GDC

March 20-26, 2022
San Francisco, CA



PHYSICS

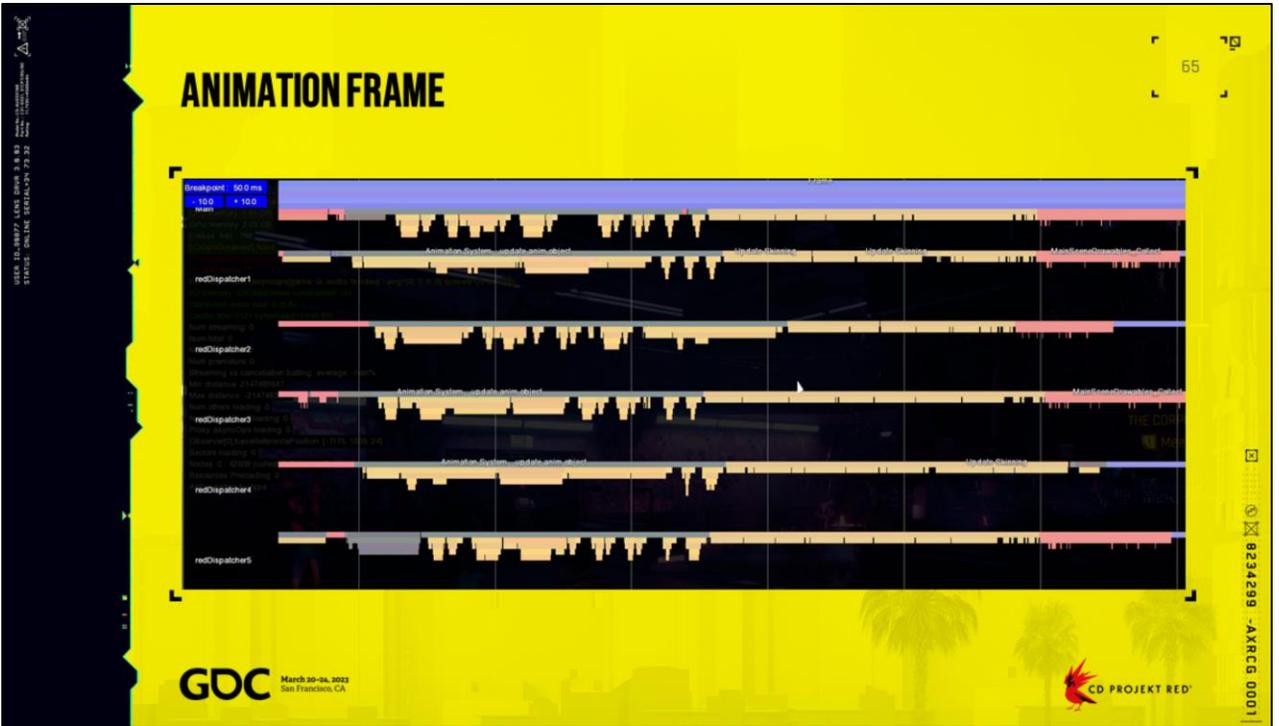
- Physics systems are built on top of PhysX
- PhysX tasks were adapted to be compatible with our job system
- Simple "C style" public API
- Safe to read and write states
- State writes are buffered and applied to all modified proxies at specific points in frame
- Concurrent writes could have been supported. However it wasn't needed

```
REDPHYSICS_API ProxyID CreateProxy( ObjectDesc& desc );  
REDPHYSICS_API void DestroyProxy( ProxyID id );  
  
REDPHYSICS_API Vector3 GetPosition( ProxyID proxyId, ActorIndex subPart = 0 );  
REDPHYSICS_API Quaternion GetRotation( ProxyID proxyId, ActorIndex subPart = 0 );  
REDPHYSICS_API Transform GetTransform( ProxyID proxyId, ActorIndex subPart = 0 );  
REDPHYSICS_API Vector3 GetLinearVelocity( ProxyID proxyId, ActorIndex subPart = 0 );  
REDPHYSICS_API Vector3 GetAngularVelocity( ProxyID proxyId, ActorIndex subPart = 0 );  
REDPHYSICS_API Float GetLinearSpeed( ProxyID proxyId, ActorIndex subPart = 0 );  
REDPHYSICS_API Vector3 GetDisplacement( ProxyID proxyId, ActorIndex subPart = 0 );  
REDPHYSICS_API ImpulseData GetImpulseAccumulator( ProxyID proxyId, ActorIndex subPart = 0 );  
REDPHYSICS_API Float GetAngularDamping( ProxyID proxyId, ActorIndex subPart = 0 );  
REDPHYSICS_API Float GetLinearDamping( ProxyID proxyId, ActorIndex subPart = 0 );
```

ANIMATION IN A NUTSHELL

64

- **Parallel update of characters**
 - As soon as character pose is calculated, schedule skinning job and send result to rendering
- **Animation instancing for massive standing crowd**
- **No update if occluded**
 - Exception if closer than 5m from player
- **Sleep mode for doors & vehicles if no movement**
- **Temporary allocation solved using scratchpad buffer**
- **Animation Streaming of 40mb budget**
 - 3k-4k animations in game



PS4 capture, a bit less than 3ms, Lizzies bar 20 npc in view, 40npc in surrounding

AI – LIVING ON THE EDGE

- Low hanging fruit & optimization of algorithms used
- Reducing behavior trees update frequency
- Parallelizing processing
- Fixing cache misses
- Logic LOD





You can change the img on the right

PS4 WITCHER 3 VS CYBERPUNK

68

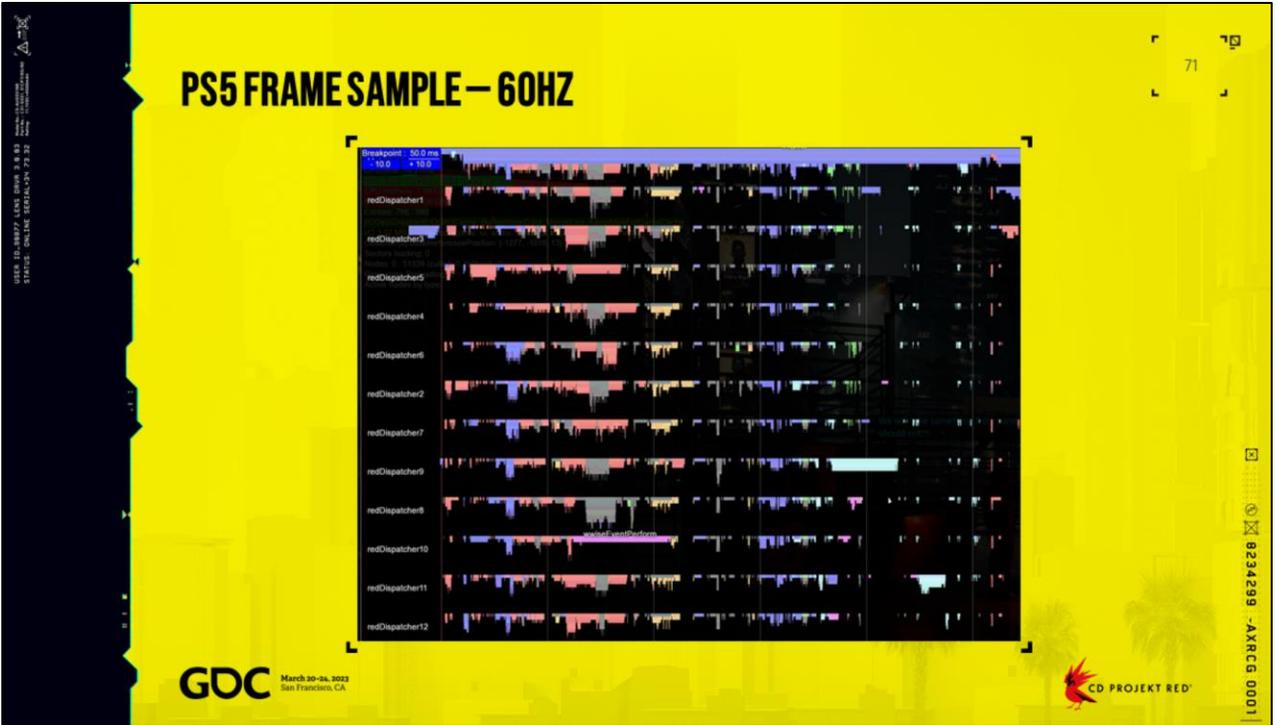


GDC March 20-24, 2022 San Francisco, CA



8234299 - AXRC6 0001

PS4, Witcher3 on top, Cyberpunk bottom one



More or less 12ms on CPU

PC FRAME SAMPLE — 6 DISPATCHERS

73

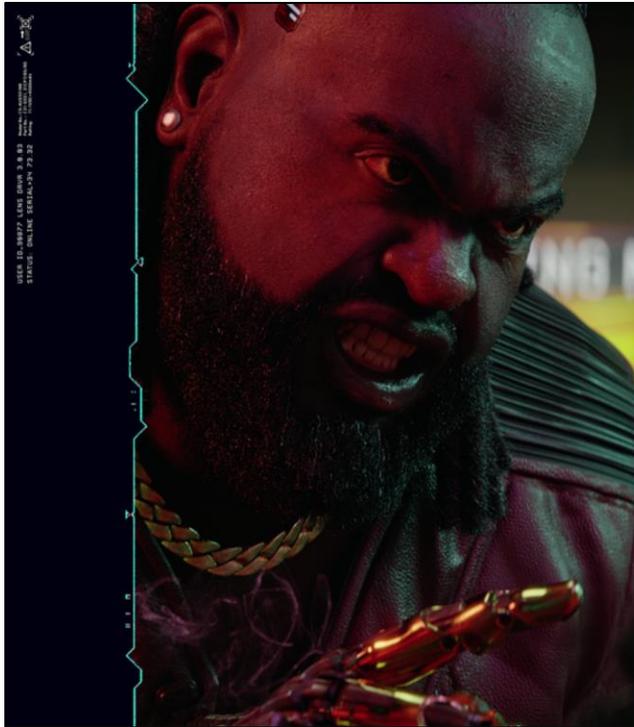
hdDispatcher1
hdDispatcher2
hdDispatcher3
hdDispatcher4
hdDispatcher5
hdDispatcher6
hdDispatcher7

GDC March 30 - 04, 2022 San Francisco, CA

CD PROJEKT RED

8234299 - AXRC6 0001

Force 6 dispatcher,
i9 - 7980XE @2.6ghz



WHAT ABOUT THE MAIN THREAD?

- We did not manage to completely eliminate it
- It only schedules beginning of frame and waits for completion
- Acts as a regular job worker while waiting

MEMBER OF THE GDC GROUP
GDC GROUP
MEMBER OF THE GDC GROUP
GDC GROUP



THANK YOU

GDC March 20-24, 2022
San Francisco, CA



8234299 - AXRCG 0001

ИЗДАТЕЛЬСТВО «СЭНТ-ПЕТЕРБУРГ»
ИЗДАТЕЛЬСТВО «СЭНТ-ПЕТЕРБУРГ»
ИЗДАТЕЛЬСТВО «СЭНТ-ПЕТЕРБУРГ»

CYBERPUNK
2-0-7-7

SAMPLES



GRAPHICS – RENDER GRAPH – SETUP & CULLING

```
auto N_StartRender = RENDER_UNIQUE_SIMPLE_COMMAND_LIST( G_None, "StartRender", UNIQUE_RENDERNODE_StartRender, CRenderNode_StartRender );
auto N_PrepareCollector = ADD_NODE( G_None, "PrepCollector", CRenderNode_PrepareCollector );
factory.Link( N_StartRender, N_PrepareCollector );

auto N_CullingScene = ADD_NODE( G_Culling, "DoCullingMainScene", CRenderNode_DoCulling, CullingMode::MainScene );
auto N_CullingRayTracing = ADD_NODE( G_Culling, "DoCullingRayTracing", CRenderNode_DoCulling, CullingMode::RayTracedObjects );
auto N_CullingCascades = ADD_NODE( G_Culling, "DoCullingCascades", CRenderNode_DoCulling, CullingMode::Cascades );
auto N_CullingLocalShadows = ADD_NODE( G_Culling, "DoCullingLocalShadows", CRenderNode_DoCulling, CullingMode::DynamicShadows );

factory.Link( N_PrepareCollector, G_Culling );
factory.Link( N_CullingScene, N_CullingLocalShadows );

// Visibility system might generate some debug geometry, so after all the culling is finished we flush it all to vertex/index data.
auto N_VisDebug = ADD_NODE( G_None, "FlushVisiDebug", CRenderNode_FlushVisibilityDebug );
factory.Link( G_Culling, N_VisDebug );

// After we've done culling that might include particles, we can kick off simulation for any particle systems that are visible.
for( UInt32 i = 0; i < NUM_PARTICLE_THREADS; ++i )
{
    ADD_NODE( G_ParticlesOnScreenSim, "SimulateOnScreenCPUParticles", CRenderNode_SimulateOnScreenCPUParticles,
        vis::CollectThreadingSetup( i, NUM_PARTICLE_THREADS ) );
}
factory.Link( N_CullingScene, G_ParticlesOnScreenSim );
factory.Link( N_CullingCascades, G_ParticlesOnScreenSim );
```

GDC

March 20-24, 2022
San Francisco, CA

CD PROJEKT RED

Set up for render, and kick off culling

"Unique" nodes are for special points in the frame. Not too important in this example, but used when merging multiple graphs -- unique stay unique (all unique from all subgraphs are de-duplicated), regular nodes are copied over.

"Simple command list" just means it creates a command list that only runs a single subnode.

"Add node" is for pure CPU work, no command list created.

Add CPU dependency, so PrepareCollector won't run until StartRender is finished.

Add nodes for doing different types of culling. Main scene (player camera), RT (inflated frustum, and area around camera), Cascades, and Local shadows (spot lights).

All culling needs to wait for PrepareCollector to finish.

In addition, LocalShadows needs to wait for main scene culling, in order to know what lights are visible.

GRAPHICS — RENDER GRAPH

```
auto N_PreRender = RENDER_COMMAND_LIST( G_None, "PreRender" )
{
    // Declare some render targets used throughout the frame. Easier than declaring them on first use, in case that first use moves.
    ADD_SUBNODE( "DeclCommonResAllocs", CRenderNode_DeclareCommonResourceAllocs );

    ADD_SUBNODE( "PrepSceneRendering", CRenderNode_PrepareSceneRendering );
    ADD_SUBNODE( "UpdateParticleRenderData", CRenderNode_UpdateParticlesRenderData );
    ADD_SUBNODE( "SpawnDynamicDecals", CRenderNode_SpawnDynamicDecals );
}
Factory.Link( G_ParticlesOnScreenSim, N_PreRender );
```

GDC March 20-26, 2022 San Francisco, CA

CD PROJEKT RED

81

0234299 - AXRC6 0001

Initial setup is ready, so we start doing some rendering work.

Command lists are explicitly defined. All subnodes in a command list node with run sequentially, although not necessarily on the same thread. Subnodes are able to branch off into additional parallel work if needed (but only for CPU work, the command list is only accessible from a single thread at a time).

In this case, UpdateParticleRenderData needs the results of the particle sim, since it's sending the final particle data to the GPU.

GRAPHICS – RENDER GRAPH – GBUFFER

82

```
auto N_GBufferPrepare = ADD_NODE( G_Render, "GBuffer_Prepare", CRenderNode_PrepareRenderElements, "renderstage_gbuffer_regular",
    SPL_OptimizedDistanceBatching, UsePreparedChunks( c_gbufferSplit ) );
factory.Link( N_CullingScene, N_GBufferPrepare );
factory.Link( G_ParticlesOnScreenSim, N_GBufferPrepare );

for( Uint32 i = 0; i < c_gbufferSplit; ++i )
{
    auto N_GBuffer = RENDER_COMMAND_LIST( G_Render, "GBuffer" )
    {
        ADD_SUBNODE( "BindGlobalConstants", CRenderNode_BindGlobalConstants );
        ADD_SUBNODE( "SetRenderToGBuffer", CRenderNode_SetRenderTargetsGBuffer, i == 0 ? rt_GBuffer_Clear : rt_GBuffer_NoClear );

        ADD_SUBNODE( "RenderElements", CRenderNode_RenderElements, "renderstage_gbuffer_regular", SPL_OptimizedDistanceBatching,
            UsePreparedChunks( c_gbufferSplit, i ) );

        ADD_SUBNODE( "EndRenderToGBuffer", CRenderNode_EndRenderTargetsGBuffer, i == 0 ? rt_GBuffer_Clear : rt_GBuffer_NoClear );
        ADD_SUBNODE( "UnbindGlobalConstants", CRenderNode_UnbindGlobalConstants );
    }

    factory.Link( N_GBufferPrepare, N_GBuffer );
}

SYNC_SUBMIT( G_None, "RenderGraphCamera_GBuffer", GpuApi::CommandListSyncType::None );
```

GDC March 20-26, 2022
San Francisco, CA

CD PROJEKT RED

GBuffer

Since we tend to have a lot of work to do in rendering the static meshes in the scene, gbuffer is split across multiple command lists.

First a CPU-only node to build and sort the list of objects to be drawn. This needs the results of main scene culling, as well as on-screen particles.

Then several command lists that each take a portion of the collected objects, drawing them to the gbuffer. These all depend on GBufferPrepare.

"BindGlobalConstants" / "UnbindGlobalConstants" are reused in many places, they set up some global constant buffers, resource bindings, etc.. Using subnodes allows that to be reused easily. Similar with setting some common render target setups.

Since we have a pretty hefty amount of work built up with the GBuffer, we might want to submit it to the GPU already, so that it can keep busy with that while we prepare more.

`SYNC_SUBMIT` will add a node that automatically has a dependency on any GPU-related nodes before it, and will submit all of them to the GPU. Here we don't need to do any additional synchronization on the GPU, so we pass `None` for `sync` type.

GRAPHICS – RENDER GRAPH – GBUFFER + VELOCITY BUFFER

83

```
auto N_GBufferVelocity = RENDER_COMMAND_LIST( G_Render, "GBufferVelocity" )
{
    ADD_SUBNODE( "BindGlobalConstants", CRenderNode_BindGlobalConstants );
    ADD_SUBNODE( "InitVelocityBuffer", CRenderNode_InitVelocityBuffer );
    ADD_SUBNODE( "BeginRenderToGBufferWithVel", CRenderNode_SetRenderTargetsGBufferWithVelocityBuffer, true );

    ADD_SUBNODE( "RenderElements", CRenderNode_RenderElements, "renderstage_gbuffer_velbuff_regular", SPL_OptimizedDistanceBatching );

    ADD_SUBNODE( "EndRenderToGBufferWithVel", CRenderNode_SetRenderTargetsGBufferWithVelocityBuffer, false );
    ADD_SUBNODE( "UnbindGlobalConstants", CRenderNode_UnbindGlobalConstants );
}
factory.Link( N_CullingScene, N_GBufferVelocity );
factory.Link( G_ParticlesOnScreenSim, N_GBufferVelocity );
```

GDC

March 20-24, 2022
San Francisco, CA



0234299 - AXRC6 0001

GBuffer + Velocity Buffer

Dynamic objects generate normal GBuffer outputs, plus they write motion vectors to a velocity buffer.

As with GBuffer, we need to wait for scene culling and particles, but it can run on the CPU in parallel with the static GBuffer.

GRAPHICS – RENDER GRAPH – ASYNC COMPUTE

84

```
SYNC_SUBMIT( G_None, "RenderGraphCamera_FinishGBuffer", GpuApi::CommandListSyncType::ForkAsyncCompute );
auto N_ASUpdateStatic = COMPUTE_COMMAND_LIST( G_Render, "AccelerationStructureUpdateStatic" )
{
    ADD_SUBNODE( "AccelerationStructureUpdateStatic", CRenderNode_AccelerationStructureUpdateStatic );
}
auto N_ASUpdateDynamic = COMPUTE_COMMAND_LIST( G_Render, "AccelerationStructuresUpdateDynamic" )
{
    ADD_SUBNODE( "AccelerationStructureUpdateDynamic", CRenderNode_AccelerationStructureUpdateDynamic );
}
auto N_ASUpdateEpilogue = COMPUTE_COMMAND_LIST( G_Render, "AccelerationStructuresUpdateEpilogue" )
{
    ADD_SUBNODE( "AccelerationStructureUpdateEpilogue", CRenderNode_AccelerationStructureUpdateEpilogue );
}

factory.Link( N_CullingRayTracing, N_ASUpdateStatic );
factory.Link( N_CullingRayTracing, N_ASUpdateDynamic );

factory.Link( N_ASUpdateStatic, N_ASUpdateEpilogue );
factory.Link( N_ASUpdateDynamic, N_ASUpdateEpilogue );

COMPUTE_COMMAND_LIST( G_Render, "AsyncComputeDuringShadowmaps" )
{
    ADD_SUBNODE( "BuildDepthChain", CRenderNode_BuildDepthChain );
    ADD_SUBNODE( "RenderSSAO", CRenderNode_DrawConeAO );
    ADD_SUBNODE( "GenerateTonemappingLUT", CRenderNode_GenerateTonemappingLUT );
    ADD_SUBNODE( "GenerateAsyncDynamicTextures", CRenderNode_GenerateAsyncDynamicTextures );
}
```

GDC

March 20-24, 2022
San Francisco, CA

CD PROJEKT RED

Async Compute work during shadowmap rendering

With the GBuffer finished, we can run some passes like a Hi-Z generation, SSAO, as well as some independent work. We run those things on async compute, while on the graphics queue we fill in shadow maps.

Our async is limited to a fork-join model, where we branch off to run specific graphics and compute work in parallel, and then sync the queues at the end. So we need to submit everything we have so far, and indicate that we want to Fork.

If we're doing ray tracing, we can build acceleration structures. Static and dynamic bottom levels have different work involved, so we split them into separate command lists. In addition those subnodes can branch off into further parallel work to prepare the acceleration structures for building. After the bottom levels are built, we have an additional step to build the top level and shader table.

We also have another compute command lists for other compute workloads, there aren't any CPU dependencies

so we don't need to link it to anything.

GRAPHICS – RENDER GRAPH – SHADOWMAPS

85

```
for( UInt32 cascadeIndex = 0; cascadeIndex < Config::GGlobalRenderingSettings.NumShadowCascades; cascadeIndex++ )
{
    auto cascadeID = RENDER_COMMAND_LIST( G_Render, "RenderCascade" )
    {
        ADD_SUBNODE( "BindGlobalConstants", CRenderNode_BindGlobalConstants );
        ADD_SUBNODE( "RenderCascade", CRenderNode_RenderShadowCascade, cascadeIndex );
        ADD_SUBNODE( "UnbindGlobalConstants", CRenderNode_UnbindGlobalConstants );
    }

    factory.Link( N_CullingCascades, cascadeID );
    factory.Link( G_ParticlesOnScreenSim, cascadeID );

    for( UInt32 lightIndex = 0; lightIndex < Config::GGlobalRenderingSettings.LocalShadowsProcessedPerFrame; lightIndex++ )
    {
        auto localShadowsID = RENDER_COMMAND_LIST( G_Render, "RenderLocalShadows" )
        {
            ADD_SUBNODE( "BindGlobalConstants", CRenderNode_BindGlobalConstants );
            ADD_SUBNODE( "RenderLocalShadows", CRenderNode_RenderLocalShadowMaps, lightIndex );
            ADD_SUBNODE( "UnbindGlobalConstants", CRenderNode_UnbindGlobalConstants );
        }

        factory.Link( N_CullingLocalShadows, localShadowsID );
        factory.Link( G_ParticlesOnScreenSim, localShadowsID );
    }

    SYNC_SUBMIT( G_None, "RenderGraphCamera_Shadows", GpuApi::CommandListSyncType::JoinAsyncCompute );
}
```

GDC

March 20-26, 2022
San Francisco, CA

CD PROJEKT RED

0234299 - AXRCG 0001

Shadowmaps

With the async compute work defined, we can do the shadowmap rendering. There's a command list for each cascade, and for each local light we plan to update. We need to wait for the appropriate culling to finish, as well as particle sim, in case there were some shadow-casting mesh particles.

And now with the compute and graphics work set up, we can submit that to the GPU, indicating that we're doing a Join sync.

GRAPHICS — RENDER GRAPH — LIGHTING & POSTFX

86

```
RENDER_COMMAND_LIST( G_Render, "Lighting" )
{
    ADD_SUBNODE( "BindGlobalConstants", CRenderNode_BindGlobalConstants );
    ADD_SUBNODE( "BindLightingGlobalConstants", CRenderNode_BindLightingGlobalConstants );

    ADD_SUBNODE( "SkyScattering", CRenderNode_RenderSkyScattering );
    ADD_SUBNODE( "VolumetricFog", CRenderNode_VolumetricFog );

    ADD_SUBNODE( "RTShadows", CRenderNode_RenderRayTracedGlobalShadow );
    ADD_SUBNODE( "RTAmbientOcclusion", CRenderNode_RenderRayTracedAmbientOcclusion );
    ADD_SUBNODE( "RTReflections", CRenderNode_RenderRayTracedReflections );
    ADD_SUBNODE( "RTFilterOutput", CRenderNode_RayTracingFilterOutput );

    ADD_SUBNODE( "LightBuffers", CRenderNode_RenderLightBuffers );
    ADD_SUBNODE( "CalculateGI", CRenderNode_CalculateGI );

    ADD_SUBNODE( "LightIntegrate", CRenderNode_RenderLightsIntegrate );

    ADD_SUBNODE( "UnbindLightingGlobalConstants", CRenderNode_UnbindLightingGlobalConstants );
    ADD_SUBNODE( "UnbindGlobalConstants", CRenderNode_UnbindGlobalConstants );
}

RENDER_COMMAND_LIST( G_Render, "PostFX" )
{
    ADD_SUBNODE( "BindGlobalConstants", CRenderNode_BindGlobalConstants );

    ADD_SUBNODE( "ApplyTAA", CRenderNode_ApplyTAA );
    ADD_SUBNODE( "ApplyDof", CRenderNode_DepthOfFieldSeparable );
    ADD_SUBNODE( "ApplyBloomAndTonemapping", CRenderNode_ApplyBloomAndTonemapping );

    ADD_SUBNODE( "UnbindGlobalConstants", CRenderNode_UnbindGlobalConstants );
}
```

GDC

March 20-24, 2022
San Francisco, CA

CD PROJEKT RED

0234299 -AXRC6 0001

Lighting, PostFX

After we have the shadowmaps and async compute joined, we can calculate our lighting, post processes, etc.

If there are no CPU dependencies, we don't need to link anything and these command lists can be built at any time.

GRAPHICS – RENDER GRAPH – DEBUG

87

```
auto debugID = RENDER_COMMAND_LIST( G_Render, "Debug" )  
{  
    ADD_SUBNODE( "BindGlobalConstants", CRenderNode_BindGlobalConstants );  
    ADD_SUBNODE( "BindLightingGlobalConstants", CRenderNode_BindLightingGlobalConstants );  
  
    ADD_SUBNODE( "ApplyDebugPreview", CRenderNode_ApplyDebugPreview );  
    ADD_SUBNODE( "RenderDebugWorld", CRenderNode_RenderDebugFragments, DebugFragments::WorldAndWorldOverlay );  
  
    ADD_SUBNODE( "UnbindLightingGlobalConstants", CRenderNode_UnbindLightingGlobalConstants );  
    ADD_SUBNODE( "UnbindGlobalConstants", CRenderNode_UnbindGlobalConstants );  
}  
factory.Link( N_VisDebug, debugID );
```

GDC March 20-26, 2022
San Francisco, CA

CD PROJEKT RED

We might have some debug geometry to draw after everything else is done.

For this, we have a dependency on VisDebug from earlier, in case it needed to add anything.

GRAPHICS — RENDER GRAPH — FINAL

```
auto NOffscreenParticles = ADD_UNIQUE( G_None, "SimulateOffScreenCPUParticles", UNIQUE_RENDERNODE_OffscreenParticlesSim,
    CRenderNode_SimulateOffScreenCPUParticles );
factory.Link( N_CullingScene, NOffscreenParticles );
factory.Link( N_CullingCascades, NOffscreenParticles );

auto NEndRender = RENDER_UNIQUE_SIMPLE_COMMAND_LIST( G_None, "EndRender", UNIQUE_RENDERNODE_EndRender, CRenderNode_EndRender );
auto NFinalFlush = ADD_UNIQUE( G_None, "FlushGPU", UNIQUE_RENDERNODE_FinalFlush, CRenderNode_Synchronize,
    GPUApi::CommandListSyncType::None, "Submit_RenderGraphCamera" );
auto NPresent = ADD_UNIQUE( G_None, "Present", UNIQUE_RENDERNODE_Present, CRenderNode_Present );
auto NEndFrame = ADD_UNIQUE( G_None, "EndFrame", UNIQUE_RENDERNODE_EndFrame, CRenderNode_EndFrame );

factory.Link( G_Render, NEndRender );
factory.Link( NOffscreenParticles, NEndRender );
factory.Link( NFinalFlush, NPresent );
factory.Link( NPresent, NEndFrame );

// A couple helpers, to fill in any missing dependencies. Adds a dependency from StartRender to anything GPU-related,
// so nothing starts filling command lists before it should. And adds a dependency from anything GPU-related to
// FinalFlush, to make sure all command lists are finished before we submit them.
factory.Link( GPUtoNextGPU( N_StartRender );
factory.Link( GPUtoPreviousGPU( N_FinalFlush );
```

GDC March 20-26, 2022
San Francisco, CA



8234299 -AXRCG 0001

Main rendering is done, so we just have some assorted finalization.

We have a separate simulation pass for offscreen particles, to keep them updating but at a throttled pace.

This is another unique node, which is only important for cases where we have multiple renders in a frame (such as when a mirror is up). We only want a single offscreen simulation, and it will wait for culling from all subgraphs so it knows what's actually offscreen everywhere.

There's a final command list with some last-minute book-keeping in it, then submit everything to the GPU that hasn't been submitted yet, and present to the screen. Everything at that point is sequential on CPU, and needs to wait for all the rendering work to complete.

GRAPHICS – RENDER NODE

89

- Render graph runs in two phases – prepare and consume
- Prepare phase records resource lifetime events
- The lifetime events in prepare / consume phases must match exactly (same order, same resource desc, etc.)
- Node declares that it requires a Command List
 - This allows to issue GPU commands on consume phase
- Job Builder can be safely used to execute parallel work, or continuation work

GDC

March 30–31, 2022
San Francisco, CA



GRAPHICS – RENDER NODE

90

```
struct RenderNode_SomeEffect : public RenderNodeBase
{
    virtual RenderNodeCommandListUsage GetCommandListUsage() const override { return RenderNodeCommandListUsage::Require; }
    virtual void Execute( const RenderNodeImplContext& rctx, job::Builder* builder ) const final override
    {
        SRenderFlowTargetDesc colorDesc = rctx.GetRegularPrecisionColorTargetDesc();
        const RenderFlowNameTag tempTex = rctx.RTTempAlloc( RENDER_NAME_TAG( "TempTarget" ), colorDesc );
        // First pass, copy main scene to an intermediate texture.
        auto tempRT = rctx.RT< RFLUF_Write >( tempTex );
        auto colorRT = rctx.RT< RFLUF_Read >( RENDER_NAME_TAG( "color" ) );

        if( rctx.IsConsumePhase() ) {
            struct CB_PARAMS_1 { Vector4 colorTint; };
            const CB_PARAMS_1 cbData; // Fill in parameters as needed
            rctx.BindTextureUAV< 0 >( tempRT );
            rctx.BindTexture< 0, eCOMPUTE >( colorRT );
            rctx.SetConstant< eCOMPUTE >( cbData );
            GetRenderer()->GetShader( SHADER_NAME( "m_postfxSomeEffect_pass1" ) )->Dispatch( ... );
        }

        // Second pass, copy it back to the main scene target.
        auto tempRT = rctx.RT< RFLUF_Read >( tempTex );
        auto colorRT = rctx.RT< RFLUF_Write >( RENDER_NAME_TAG( "color" ) );
        if( rctx.IsConsumePhase() ) {
            struct CB_PARAMS_2 { Matrix uvTransform; };
            const CB_PARAMS_2 cbData; // Fill in parameters as needed
            rctx.BindTextureUAV< 0 >( colorRT );
            rctx.BindTexture< 0, eCOMPUTE >( tempRT );
            rctx.SetConstant< eCOMPUTE >( cbData );
            GetRenderer()->GetShader( SHADER_NAME( "m_postfxSomeEffect_pass2" ) )->Dispatch( ... );
        }
    }
};
```

GDC

March 20-24, 2022
San Francisco, CA

CD PROJEKT RED

0234299 - AXRC6 0001

First example, some sort of post process

This is using a temporary texture to hold some intermediate results.

The graph is run in two phases, Prepare and Consume. During Prepare phase, resource lifetime events are recorded, so that before running the Consume phase we can analyze total lifetime, resource sizes, etc. and alias multiple resources over the same chunks of GPU memory.

The lifetime events in Prepare and Consume must match exactly (same order, same resource desc's), so occur outside of any checks for the current phase.

RTTempAlloc doesn't necessarily cause the resource to be allocated, it will be available only after the first use (in this case, marked by `rctx.RT<>`). We specify how the resource is intended to be used (read or write), but this metadata ended up not really being used.

The node declares that it requires a Command List, which allows it to issue GPU commands during the Consume phase.

GRAPHICS – RENDER NODE

91

```
struct RenderNode_DoCulling : public RenderNodeBase
{
    virtual RenderNodeCommandListUsage GetCommandListUsage() const override { return RenderNodeCommandListUsage::None; }
    virtual Bool GetJobBuilderUsage() const override { return true; }
    virtual void Execute( const RenderNodeImplContext& rctx, job::Builder* builder ) const override
    {
        if( rctx.IsConsumePhase() ) {
            if( m_cullingMode == CullingMode::MainScene ) {
                // issue a job as part of this node. The node will not be finished until all additional work is done.
                builder->DispatchJob( "DoCull_SceneLayer", [rctx]( const job::RunContext& context ) {
                    // Do culling, which may spawn additional jobs as children of this one.
                } );
            }
            // etc.
        }
        CullingMode m_cullingMode;
    }
};
```

GDC March 20-26, 2022
San Francisco, CA



Next example is a CPU-only node. Because it declares that it won't use a Command List, it would be invalid to try to issue any GPU commands during its execution.

In this case, we also declare that we would like to use a `job::Builder`, which allows for spawning additional parallel work. Any dependencies on this node will need to wait until any additional jobs are also finished.

Non-GPU node, which spawns additional jobs as part of the work.

GRAPHICS – RENDER NODE

92

```
struct RenderNode_AccelerationStructureUpdateEpilogue : public RenderNodeBase
{
    virtual RenderNodeCommandListUsage GetCommandListUsage() const override { return RenderNodeCommandListUsage::Require; }
    virtual Bool GetJobBuilderUsage() const final override { return true; }
    virtual void Execute( const RenderNodeImplContext& rctx, job::Builder* builder ) const final override
    {
        if( rctx.IsConsumePhase() )
        {
            // Issue some GPU commands first

            // Unbind the command list from the current thread. Another thread could run on this thread!
            GpuApi::BindCommandList( {} );
            // Spawn several jobs to do some processing in parallel over all instances.
            CollectTLAS_Parallel( rctx, *builder );

            // After the above jobs are finished, run another job which will finish up with some more GPU work.
            builder->DispatchJob( "Finalize", [ this, rctx ]( const job::RunContext& context ) {
                // Bind this node's command list to whatever thread we're running on.
                GpuApi::BindCommandList( rctx.GetCommandList() );
                UpdateDataGPU( rctx );

                // Unbind it, because the remaining nodes in this command list can end up on a different thread.
                GpuApi::BindCommandList( {} );
            } );
        }
    }
};
```

GDC

March 20-26, 2022
San Francisco, CA

CD PROJEKT RED

8234299 - AXRC6 0001

Finally we have a GPU node, which also can spawn additional parallel work.

We only have a single Command List, so we need to be careful when accessing it so multiple threads don't try to access it. But there's no restriction on which thread it can be used.

A Command List must be bound to a thread, and can only be bound to a single thread at a time. So in order to access it from a separate job, we need to unbind it from the current thread first.

PHYSICS— STATE BUFFERING

- **Predefined states:**
 - Position, Rotation, Linear Velocity, Angular Velocity... and more!
- **State Block**
 - All information to set a state in a proxy
 - Up to 4 states in a single block
 - If more are needed, linked list of StateBlocks
- **State Allocator**
 - Can allocate a block (max 64k)
 - Can allocate data for states (preallocated 1mb)
 - In both cases, a linear allocator flushed every frame
- **State Flush**
 - At given point(s) in frame, all valid and buffered states are applied to proxies

PHYSICS—STATE BLOCK

94

All write operations are wrapped in ProxyWriter interface

1. Takes ProxyID as argument
2. Allocates new block or get current one if already exists
3. If there is an existing block, create full copy of all states so far
4. Writes all needed states to block, or allocate new one if there's no space
5. Commit states. If pointer from the beginning didn't change, swap pointers. Otherwise go back to 3 *

```
struct StateBlock
{
    enum { MAX_STATES = 4 };
    ProxyID proxyID;
    uint16_t numStates;

    union {
        uint8_t all;
        struct {
            uint8_t headOfList : 1;
            uint8_t flushed : 1;
            uint8_t outdated : 1;
        };
    };
    flags;

    StateValueFlag_ stateValueFlags;
    StateBlock* next;
    uint32_t frameNo;

    struct {
        struct Offset {
            uint32_t offset : 24;
            uint32_t size : 8;
        };

        Offset offset [MAX_STATES];
        uint8_t state [MAX_STATES];
        uint8_t subpart [MAX_STATES];
        uint8_t nbShapes [MAX_STATES];
    } states;
};
```

GDC

March 20-26, 2022
San Francisco, CA



0234299 - AXRC6 0001

Copy of all states can be done on multiple threads

* In practice there was no concurrent writers. We just swap pointers on commit. But other threads can safely read states at same time.

PHYSICS— STATE ALLOCATOR & FLUSH

95

- **StateAllocator**
 - Acts like a single linear allocator. Allocated from fixed, preallocated memory area
 - All content is discarded after flush
- **State Flush**
 - For each proxy, check if there is a StateBlock allocated. Should be marked as "headOfList"
 - Traverse through all states and connected blocks and apply data to proxy*

```
void FlushBufferedProxyStates( StateAllocator* stateAllocator ) {  
    StateBlock* blockIt = stateAllocator->BlockBegin();  
    StateBlock* blockEnd = stateAllocator->BlockEnd();  
    while (blockIt != blockEnd) {  
        if( StateBlock* proxyCurrentBlock = g_proxyLookup->m_stateBlocks[blockIt->proxyId.Index()].GetValue() ) {  
            proxy_internal::FlushProxyState( stateAllocator, proxyCurrentBlock->proxyId, proxyCurrentBlock );  
            proxy_internal::CommitStateBlock( proxyCurrentBlock->proxyId, nullptr );  
        }  
        ++blockIt;  
    }  
    stateAllocator->Clear();  
}
```

GDC

March 20-24, 2022
San Francisco, CA

CD PROJEKT RED

* This operation is single threaded. All writes to physics scene are not thread-safe.