



Structured-Light Based Acquisition (Part 2)

CS635 Spring 2010

Daniel G. Aliaga
Department of Computer Science
Purdue University

Acquiring Dynamic Scenes




- Scene: object (or camera) is moving and/or object is deforming
- Acquisition: capture as much information as possible in one to a few frames
 - By exploiting coherence
 - By exploiting several “channels” of information (e.g., color, infrared, etc...)

Acquiring Dynamic Scenes





- “Capturing 2½D Depth and Texture of Time-Varying Scenes Using Structured Infrared Light”, Frueh and Zakhor, PROCAMS 2005
 - Use single-frames and infrared illumination...
- “Rapid Shape Acquisition Using Color Structured Light and Multi-pass Dynamic Programming”, Zhang et al., 3DPVT 2002
 - Use single-frames and colored patterns...
- “Fast 3D Scanning with Automatic Motion Compensation”, Weise et al., CVPR 2007
 - Use phase shifting and motion compensation over a few frames...
- “Real-time 3D Model Acquisition”, Rusinkiewicz et al., SIGGRAPH 2002
 - Use patterns producing local correspondences over a few frames and merge...

Capturing 2½D Depth and Texture of Time-Varying Scenes Using Structured Infrared Light





V-lines




Line defined
at "middle" of
IR strip

How do you know which line is which? Ideas?



H-lines



H-line sweeps up/down at 2Hz and enables an ordering of (a subset of) the V-lines and thus permits their correspondence


Figure 6: Reconstructing the depth along V-lines. (a) IR frame; (b) V-lines from intra-frame tracking only; (c) V-lines with additional forward inter-frame tracking, (d) final result after V-lines with both forward and backward inter-frame tracking, and line counting.



Additional Steps



Grab color image



Foreground segmentation, and dense depth interpolation



Put it all together...


IR camera at 30Hz, color camera at 10Hz
(probably faster today...)

Rapid Shape Acquisition Using Color Structured Light and Multi-pass Dynamic Programming




- Recall: how do we correspond lines?

Rapid Shape Acquisition Using Color Structured Light and Multi-pass Dynamic Programming




- Use color transitions to define features
- Define lines at the transitions from color A to color B

Rapid Shape Acquisition Using Color Structured Light and Multi-pass Dynamic Programming



- What is a notable problem?
- Resolution. Why?

Rapid Shape Acquisition Using Color Structured Light and Multi-pass Dynamic Programming



- Only have three color channels (R,G,B) and can only robustly differentiate “strong” color changes
- This reduces the number of colors to use, and
- Often results in ambiguity in the color coding

Rapid Shape Acquisition Using Color Structured Light and Multi-pass Dynamic Programming



- Challenges

- Given a color code, how to do “best” correspond the stripes?
- With the above in mind, how do we design a good color code?

Rapid Shape Acquisition Using Color Structured Light and Multi-pass Dynamic Programming



- Challenges

- Given a color code, how to do “best” correspond the stripes?
- With the above in mind, how do we design a good color code?




How to “best” correspond the stripes



- Solution
 - Dynamic Programming



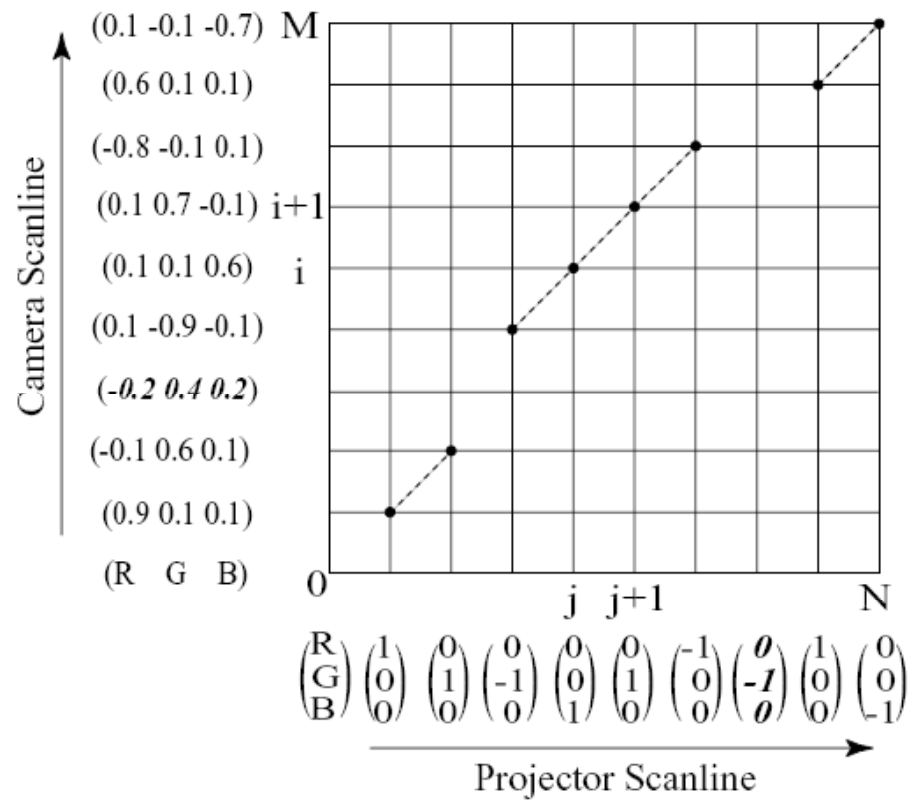



How to “best” correspond the stripes?



(rectified images)



How to “best” correspond the stripes?



How to “best” correspond the stripes?



Multiple match hypotheses $\phi = \left\{ \begin{pmatrix} j_1 \\ i_1 \end{pmatrix}, \begin{pmatrix} j_2 \\ i_2 \end{pmatrix}, \dots, \begin{pmatrix} j_H \\ i_H \end{pmatrix} \right\}$

Similarity score (of color) between
edge e_i and transition q_j is $s(q_j, e_i)$

Score of the entire match sequence $f(\phi) = \sum_{k=1}^H s(q_{j_k}, e_{i_k})$

Dynamic programming objective is: $\arg \max_{\phi} (f(\phi))$

How to “best” correspond the stripes?



Dynamic programming objective is: $\arg \max_{\phi} (f(\phi))$

However, the space all possible ϕ is very large: $O(M^N)$

Solution?

Assume monotonicity (of the depth ordering):

$$i_1 \leq i_2 \leq \dots \leq i_H$$

Great! But this monotonicity does **not** hold in what situation?

Occlusions! Oh well...

But it holds for individual fragments, which we can combine

How to “best” correspond the stripes?



Dynamic programming objective is: $\arg \max_{\phi} (f(\phi))$

Let optimal ϕ be called ϕ^*

$$f(\phi^*_{ji}) = \begin{cases} 0 & \text{if } j=0 \text{ or } i=0 \\ \max \begin{cases} f(\phi^*_{j-1,i-1}) + s(q_j, e_i) \\ f(\phi^*_{j-1,i}) \\ f(\phi^*_{j,i-1}) \end{cases} & \text{otherwise} \end{cases}$$

f found through a recursive search and some optimizations to further reduce the search space (e.g., assume at most small depth changes from one column to another)

Rapid Shape Acquisition Using Color Structured Light and Multi-pass Dynamic Programming



- Challenges

- Given a color code, how to do “best” correspond the stripes?
- With the above in mind, how do we design a good color code?

How do we design a good color code?




- De Bruijn sequence $B(k, n)$
 - is a cyclic sequence of a given alphabet A with size k for which every possible subsequence of length n in A appears as a sequence of consecutive characters exactly once
- $B(k, n)$ has length k^n
- Example: $A = \{0, 1\}$
 - $B(2, 3) = 00010111$ or 11101000

De Bruijn sequence $B(k,n)$



- Can also be constructed by taking a Hamiltonian path of an n -dimensional De Bruijn graph over k symbols; e.g.,

(Hamiltonian path means each vertex is visited once)







Color Sequence

- Colors = $\{000, 100, 110, \dots, 111\}$ total of $8-1=7$ because 000 is useless
- Color sequence is created by $p_{j+1} = p_j \text{ XOR } d_j$
 - XOR'ing effectively “flips bits” using d_j
 - p_0 is a chosen initial color (e.g., 100)
- Want 3 letters sequences d_j to be unique
- In practice about 125 stripes is sufficient
- Thus, a $B(5,3)$ is adequate



Examples



Fast 3D Scanning with Automatic Motion Compensation




Figure 1. 3D reconstructions of a static (left) and a moving (right) hand. Motion compensation (bottom right) removes the ripples from the reconstructed surface (top right).

- Higher resolution/quality than previous method
- Uses phase-shifting and motion-compensation

Fast 3D Scanning with Automatic Motion Compensation




Figure 7. Reconstruction of a complex scene containing several objects (phone, teapot, figure, fruit): a) texture image, b) reconstructed phase, c) geometry, d) textured geometry, e)+f) close-ups



Figure 8. Reconstruction of a waving cloth. Motion correction correctly removes the ripples (right).



Figure 9. Reconstruction of a person speaking.


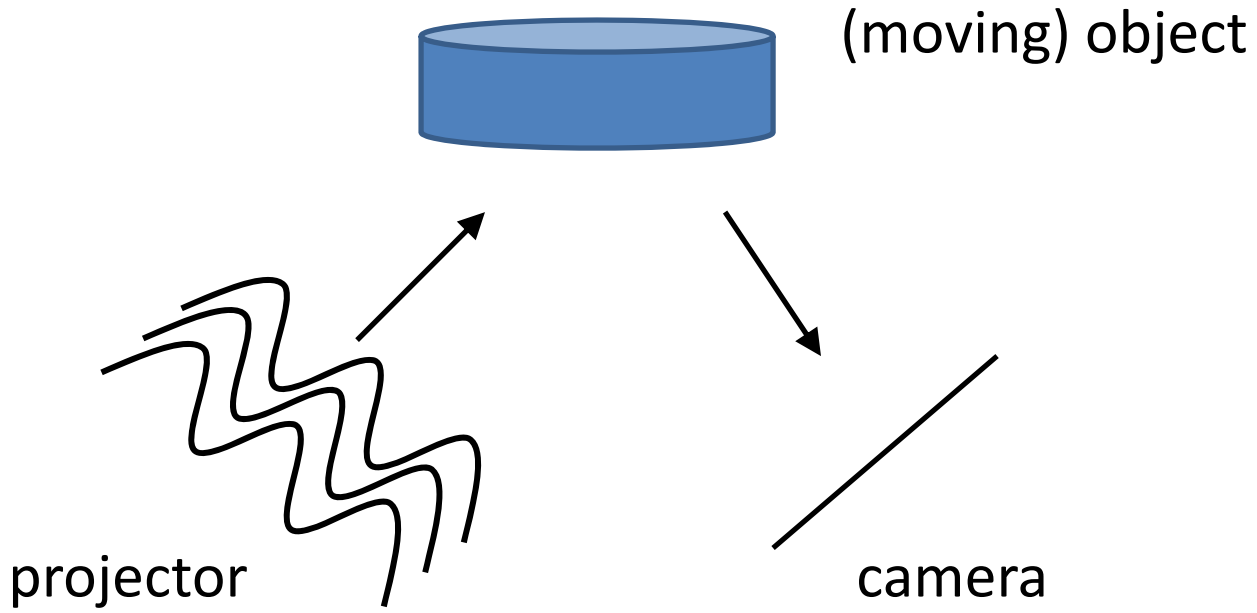


Figure 10. Reconstruction of moving hands in front of the torso. On the right with motion compensation.



Figure 11. Online reconstruction of hand gestures.

What is (standard) phase shifting?



Motion Compensation




- Since phase shifting assumes a static scene, **correlation-based stereo** is used to compensate for motion
- An additional modification is proposed to **handle discontinuities** (which also plague standard phase shifting)

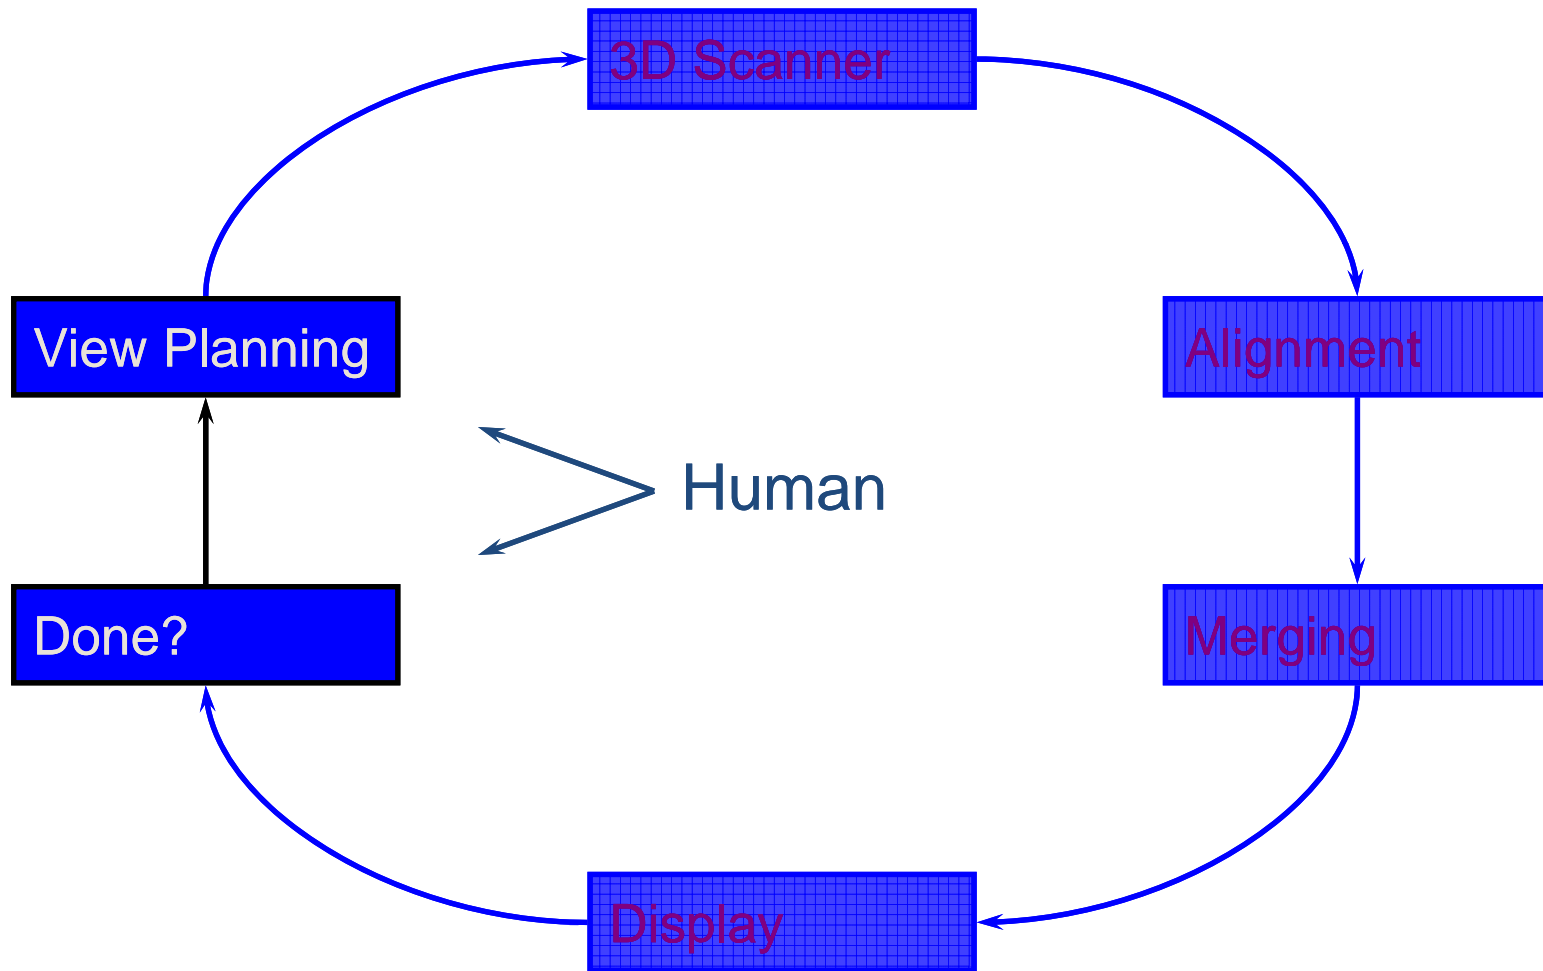
Real-Time 3D Model Acquisition



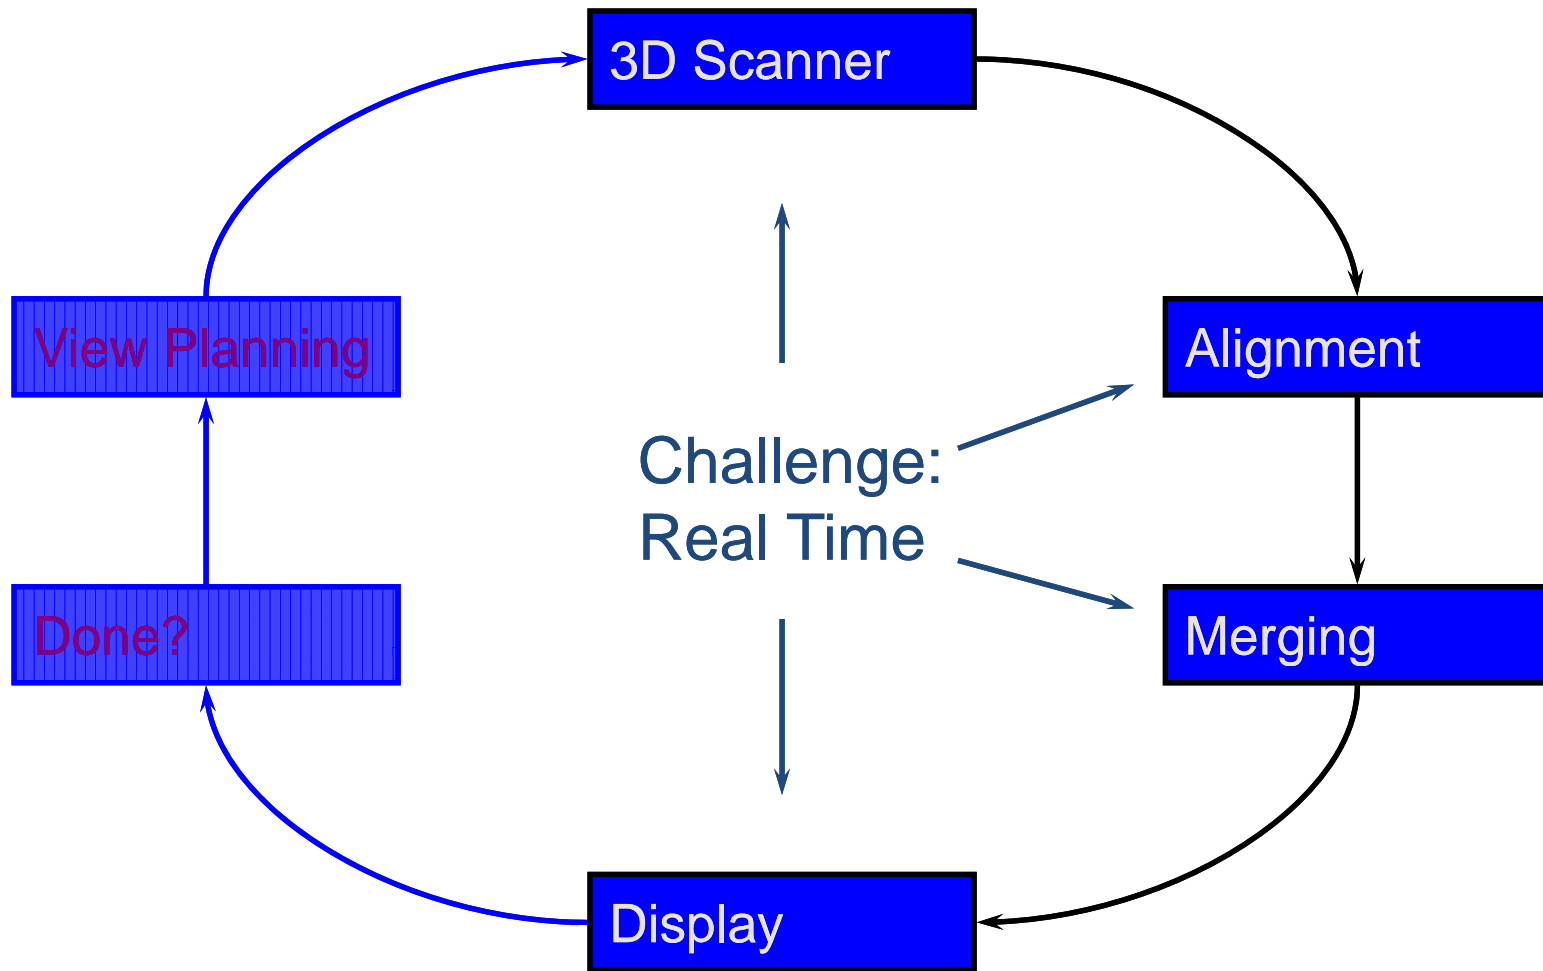
(slides and videos of this section by Szymon Rusinkiewicz @ Princeton)



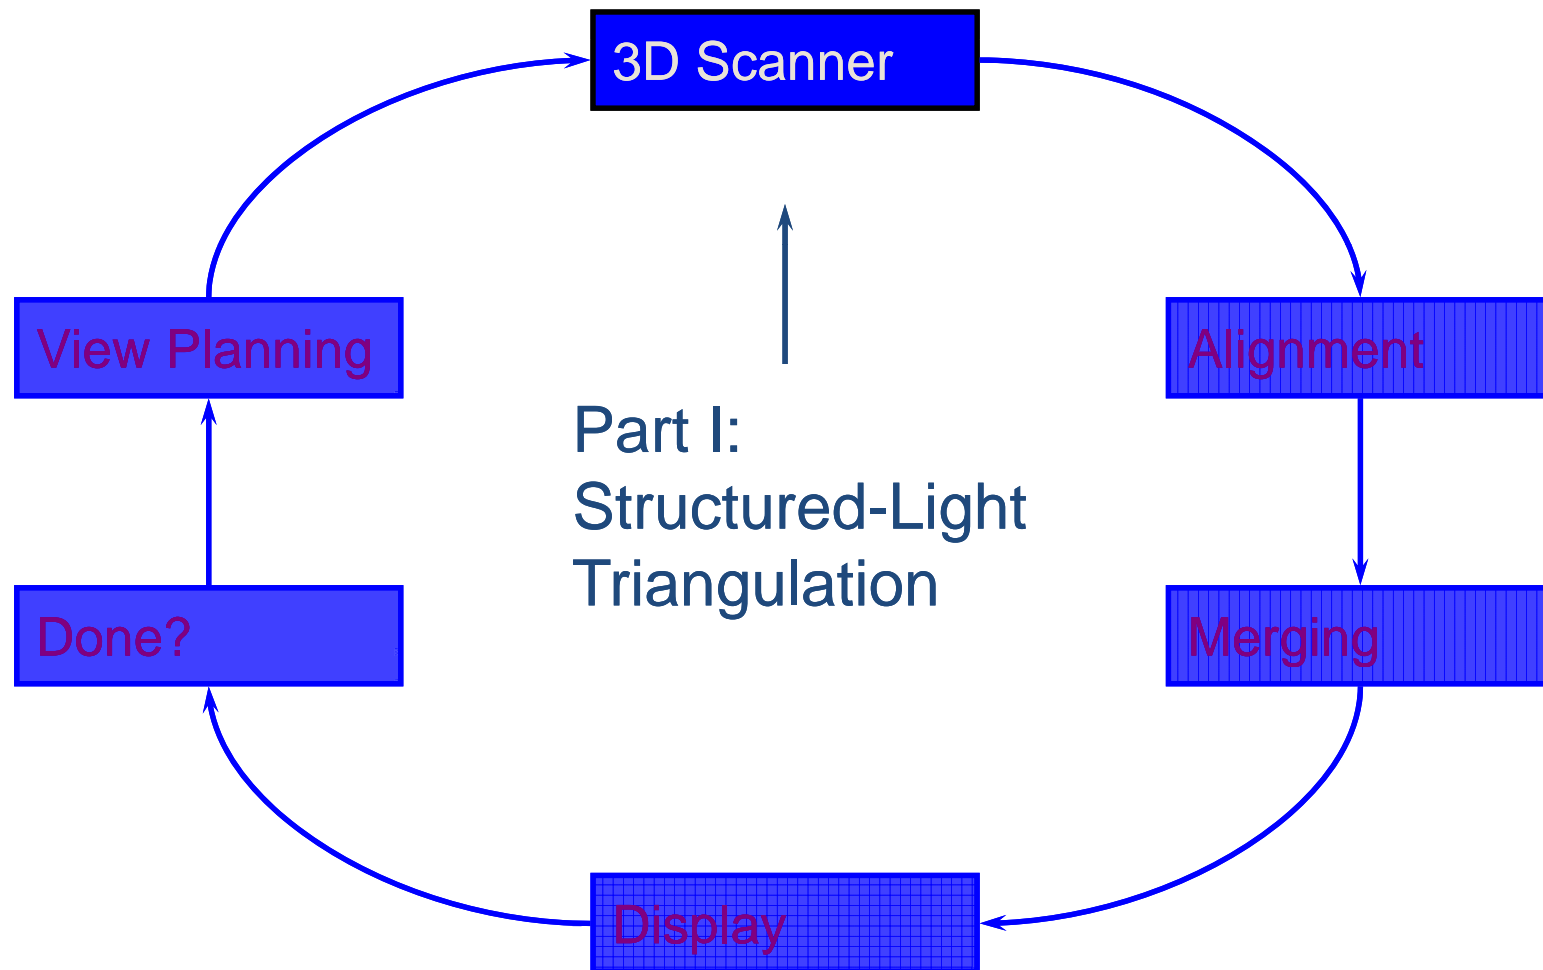
Real-Time 3D Model Acquisition Pipeline



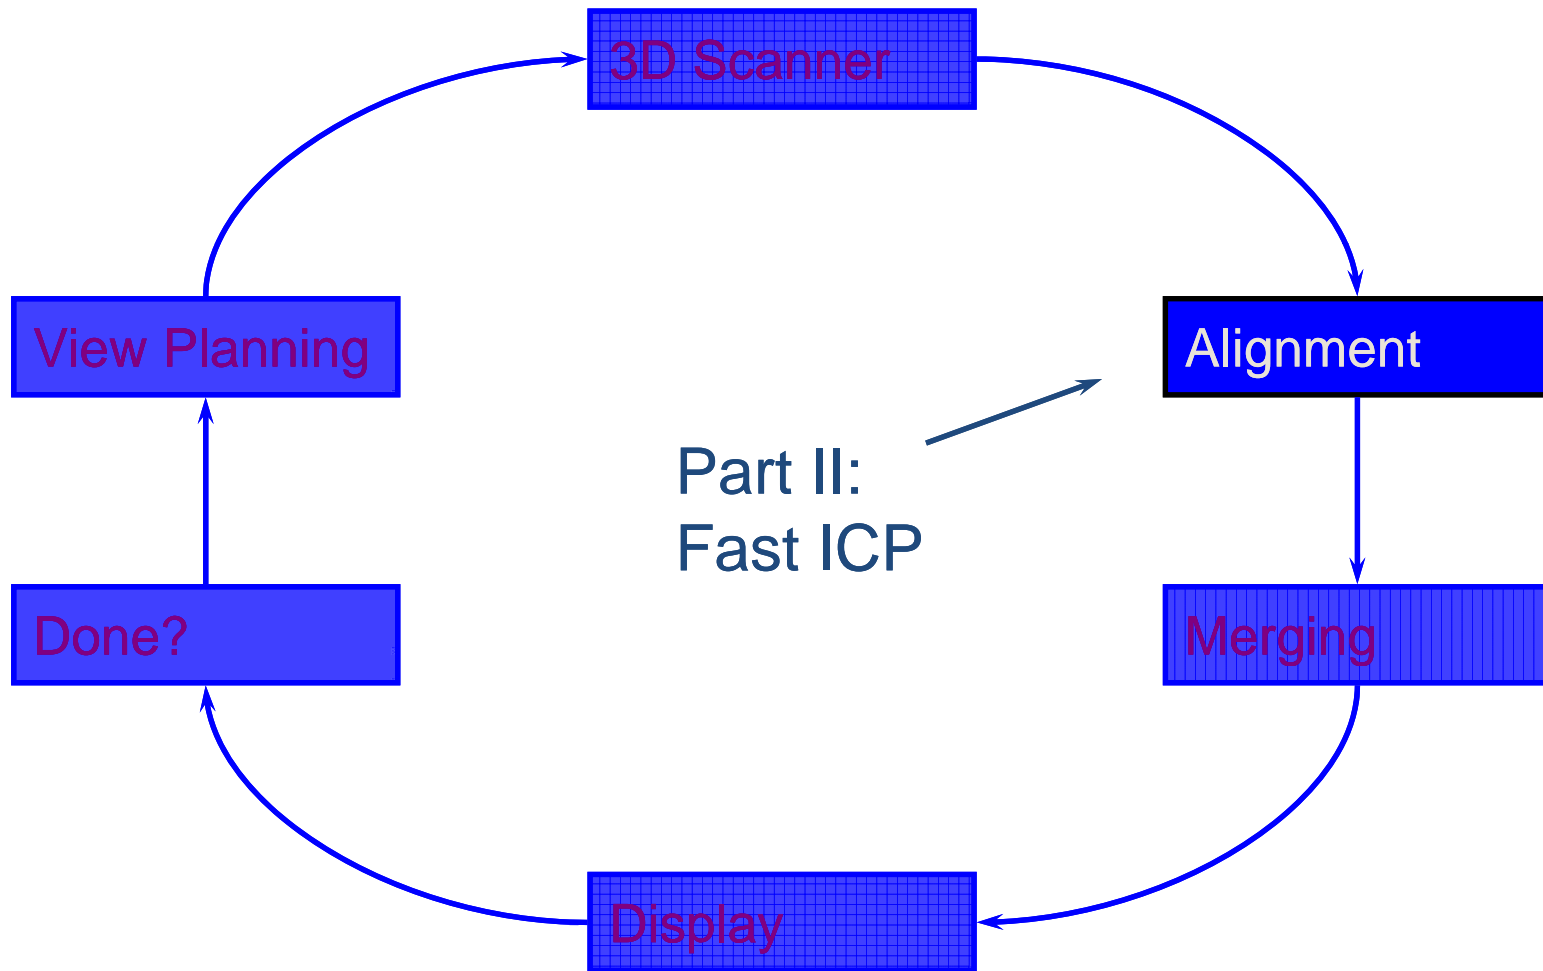
Real-Time 3D Model Acquisition Pipeline



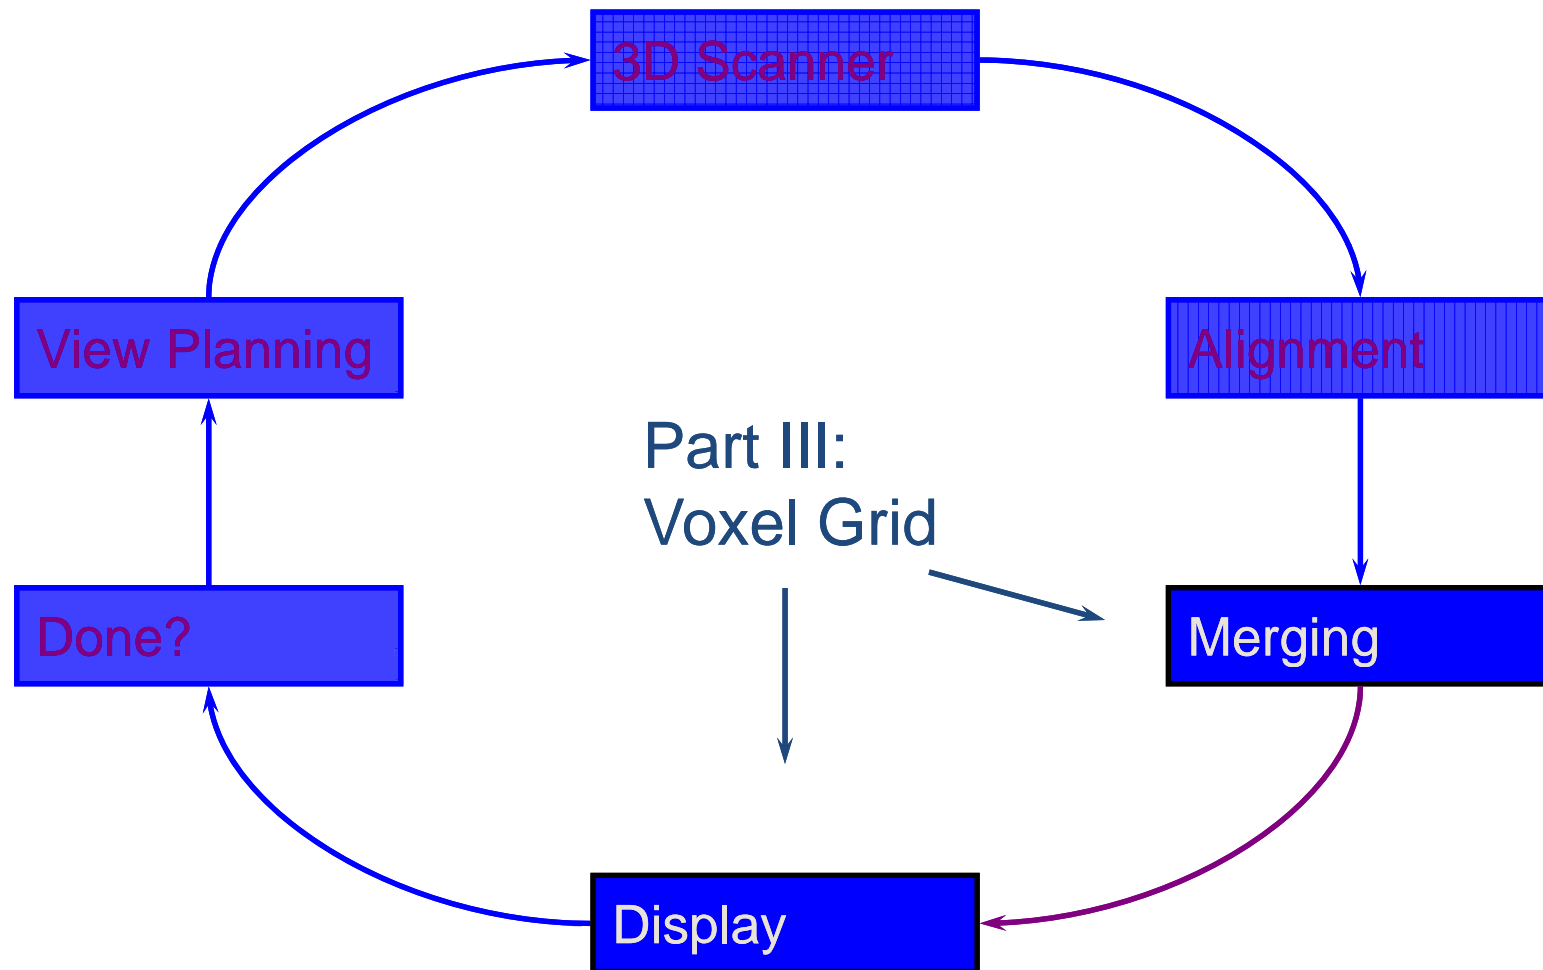
Real-Time 3D Model Acquisition Pipeline



Real-Time 3D Model Acquisition Pipeline




Real-Time 3D Model Acquisition Pipeline





Triangulation



- Depth from ray-plane triangulation



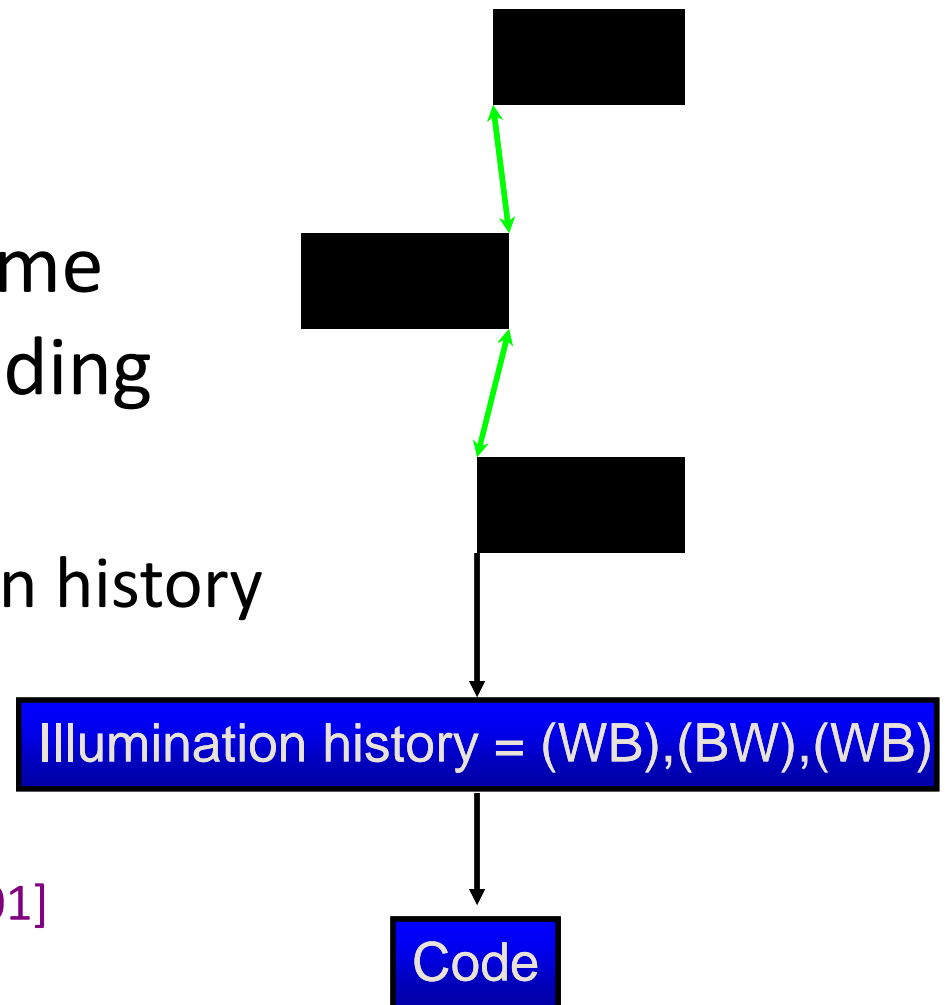
Triangulation

- Faster acquisition: project multiple stripes
- Correspondence problem: which stripe is which?

Codes for Moving Scenes



- Assign time codes to stripe boundaries
- Perform frame-to-frame tracking of corresponding boundaries
 - Propagate illumination history




[Hall-Holt & Rusinkiewicz, ICCV 2001]




Designing a Code

- Want many “features” to track:
lots of black/white edges at each frame
- Try to minimize ghosts – WW or BB
“boundaries” that can’t be seen directly










Designing a Code



[Hall-Holt & Rusinkiewicz, ICCV 2001]


Space-Time Boundary Code






Implementation

- Pipeline:



- DLP projector illuminates scene @ 60 Hz.
- Synchronized NTSC camera captures video
- Pipeline returns range images @ 60 Hz.

Real-Time 3D Model Acquisition Pipeline






Aligning 3D Data

- This range scanner can be used for any moving objects
- For **rigid objects**, range images can be aligned to each other as object moves



Aligning 3D Data


- ICP (Iterative Closest Points): for each point on one scan, minimize distance to closest point on other scan...





Aligning 3D Data

- ... and iterate to find alignment
 - Iterated Closest Points (ICP) [Besl & McKay 92]





ICP in the Real-Time Pipeline



- Potential problem with ICP: local minima
 - In this pipeline, scans close together
 - Very likely to converge to correct (global) minimum
- Basic ICP algorithm too slow (~ seconds)
 - Point-to-plane minimization
 - Projection-based matching
 - With these tweaks, running time ~ milliseconds

[Rusinkiewicz & Levoy, 3DIM 2001]

Real-Time 3D Model Acquisition Pipeline




Merging and Rendering




- Goal: visualize the model well enough to be able to see holes
- Cannot display all the scanned data – accumulates linearly with time
- Standard high-quality merging methods: processing time \sim 1 minute per scan


Merging and Rendering




Merging and Rendering




Merging and Rendering



Merging and Rendering



Merging and Rendering




- Point rendering, using accumulated normals for lighting

Example: Photograph



18 cm.

Result






Postprocessing

- Real-time display
 - Quality/speed tradeoff
 - Goal: let user evaluate coverage, fill holes
- Offline postprocessing for high-quality models
 - Global registration
 - High-quality merging (e.g., using VRIP [Curless 96])

Postprocessed Model



Recapturing Alignment





Summary

- 3D model acquisition pipeline optimized for obtaining complete, hole-free models
- Use human's time most efficiently
- Pieces of pipeline selected for real-time use:
 - Structured-light scanner for moving objects
 - Fast ICP variant
 - Simple grid-based merging, point rendering



Limitations

- Prototype noisier than commercial systems
 - Could be made equivalent with careful engineering
 - Ultimate limitations on quality: focus, texture
- Scan-to-scan ICP not perfect \Rightarrow alignment drift
 - Due to noise, miscalibration, degenerate geometry
 - Reduced, but not eliminated, by “anchor scans”
 - Possibly combine ICP with separate trackers