

Compressing Dynamic Data Structures*

Haifeng He (hehf@cs.arizona.edu)

Computer Science
Department

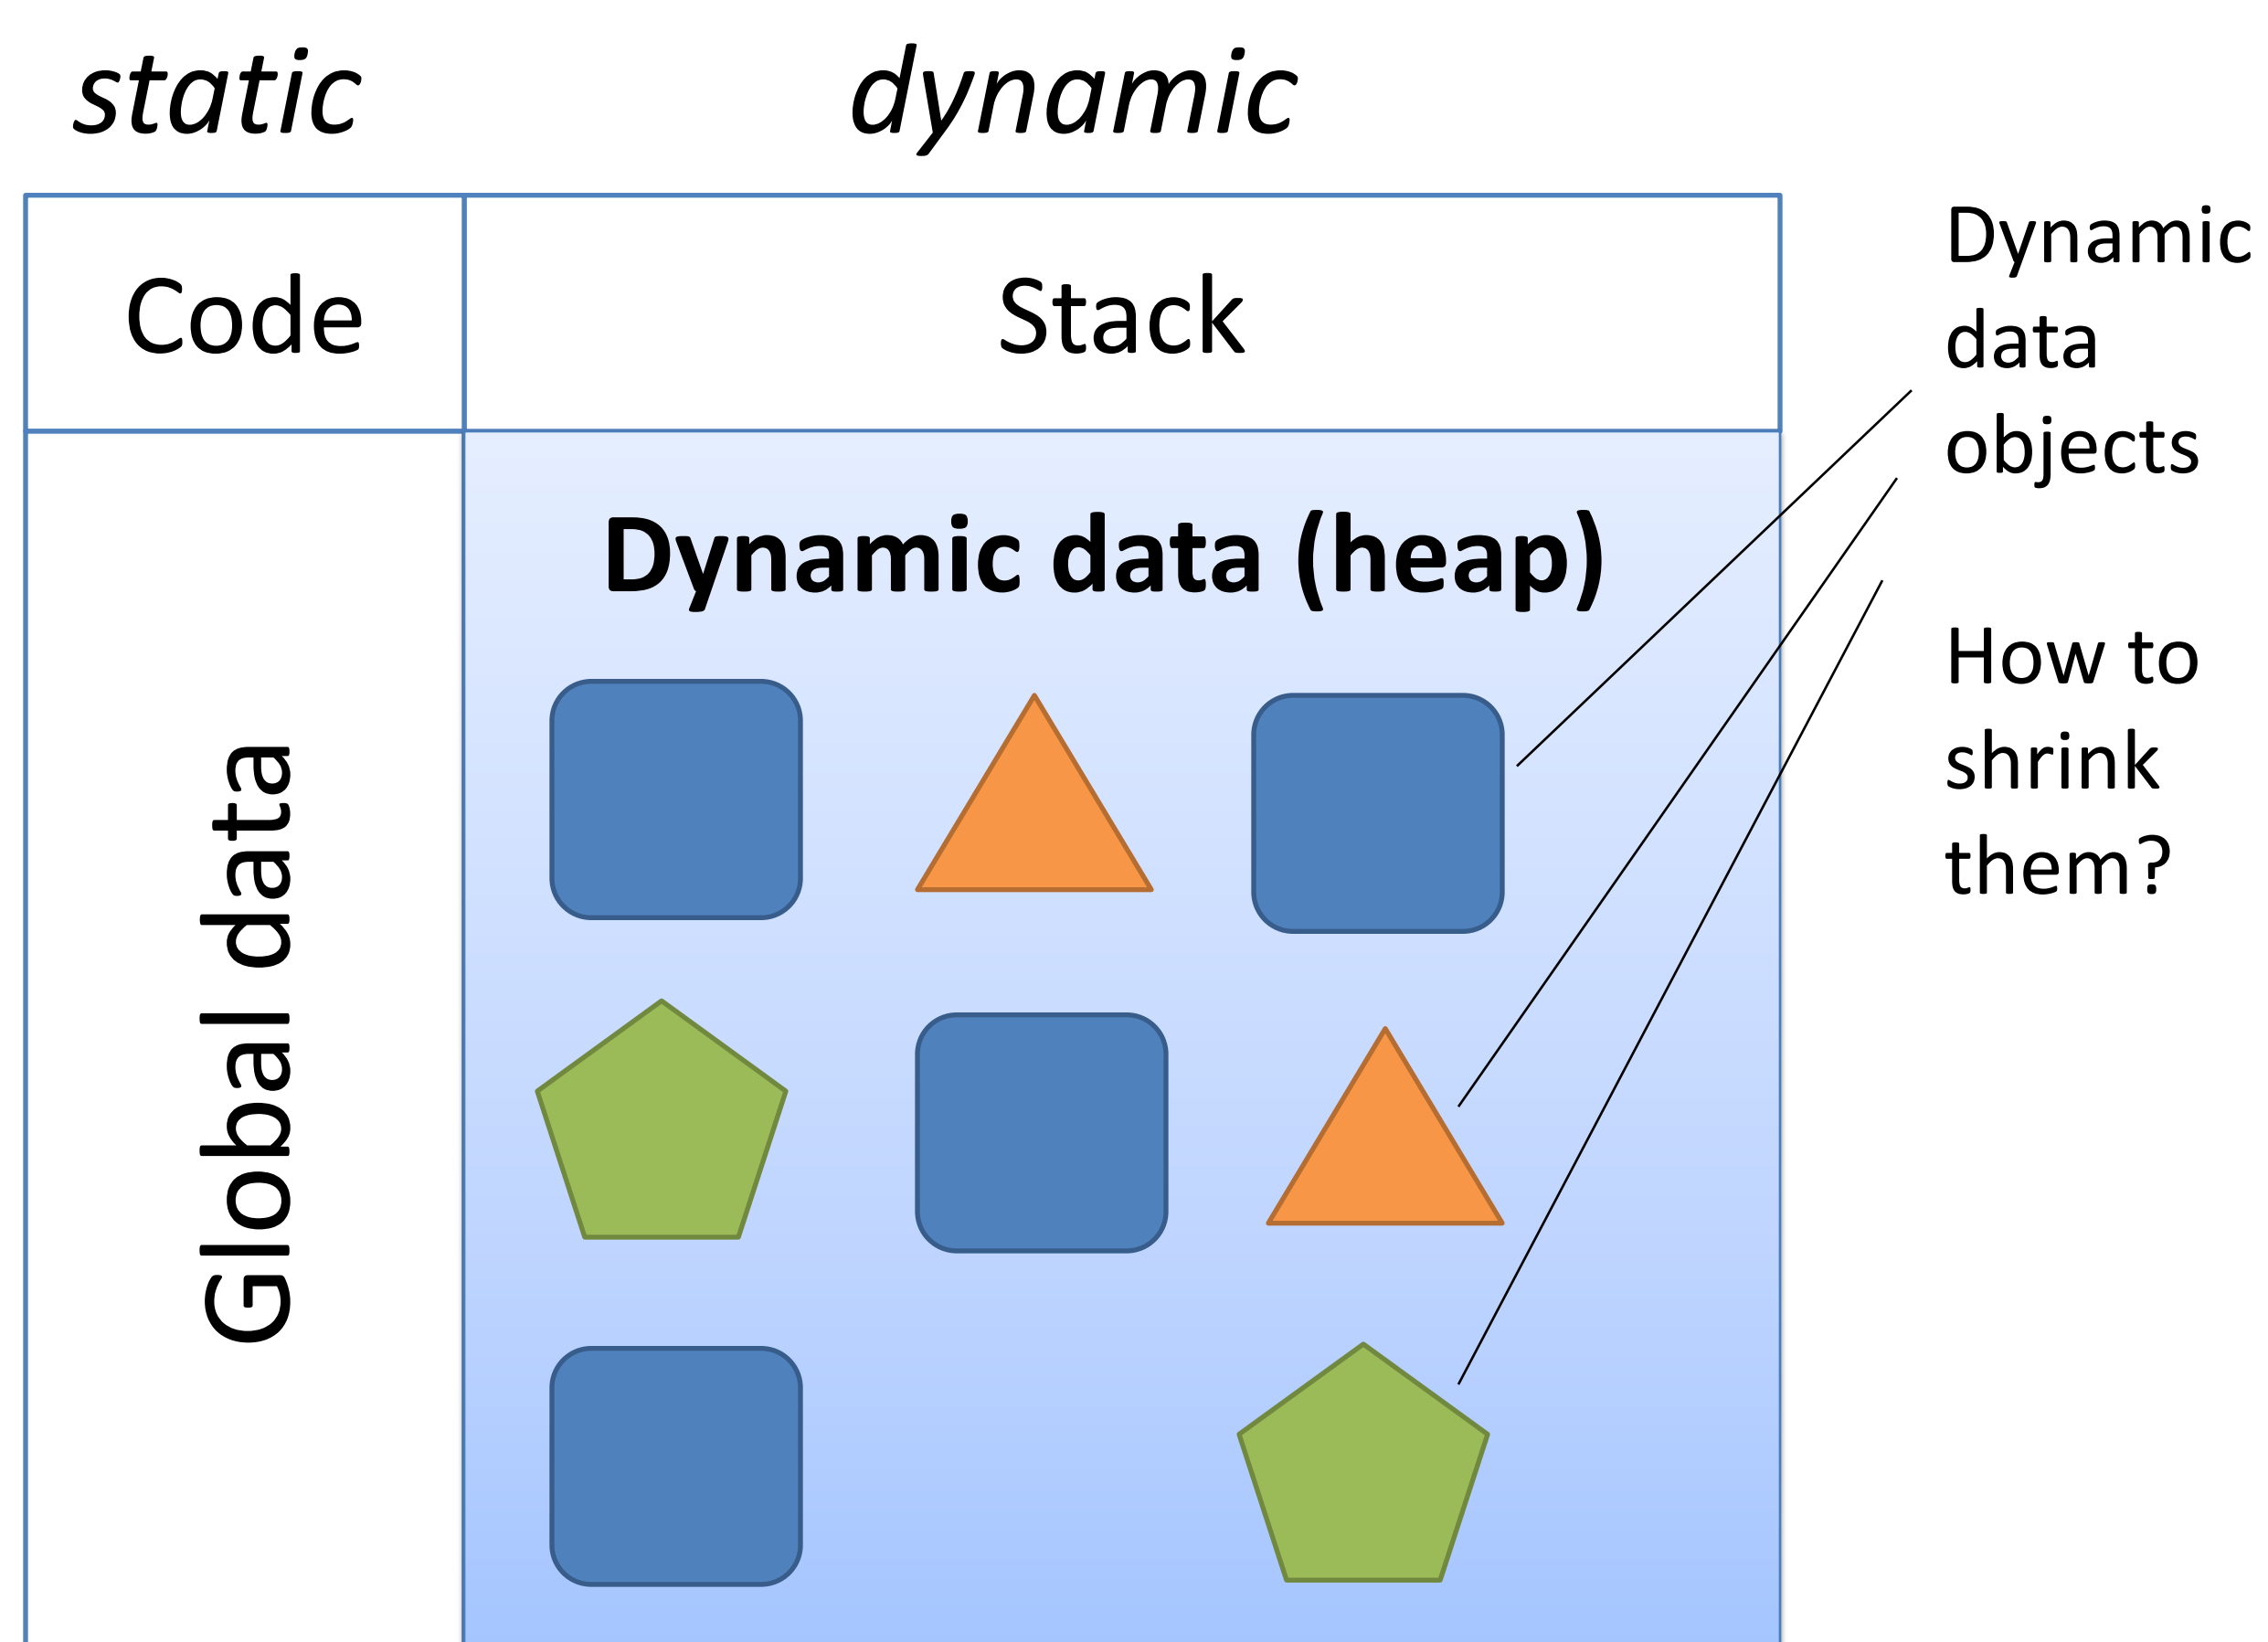


Problem



- Embedded system software is becoming increasingly complex.
- Physical memory is limited.
- Reduce memory consumption by *compressing dynamic data structures*

Program memory footprint

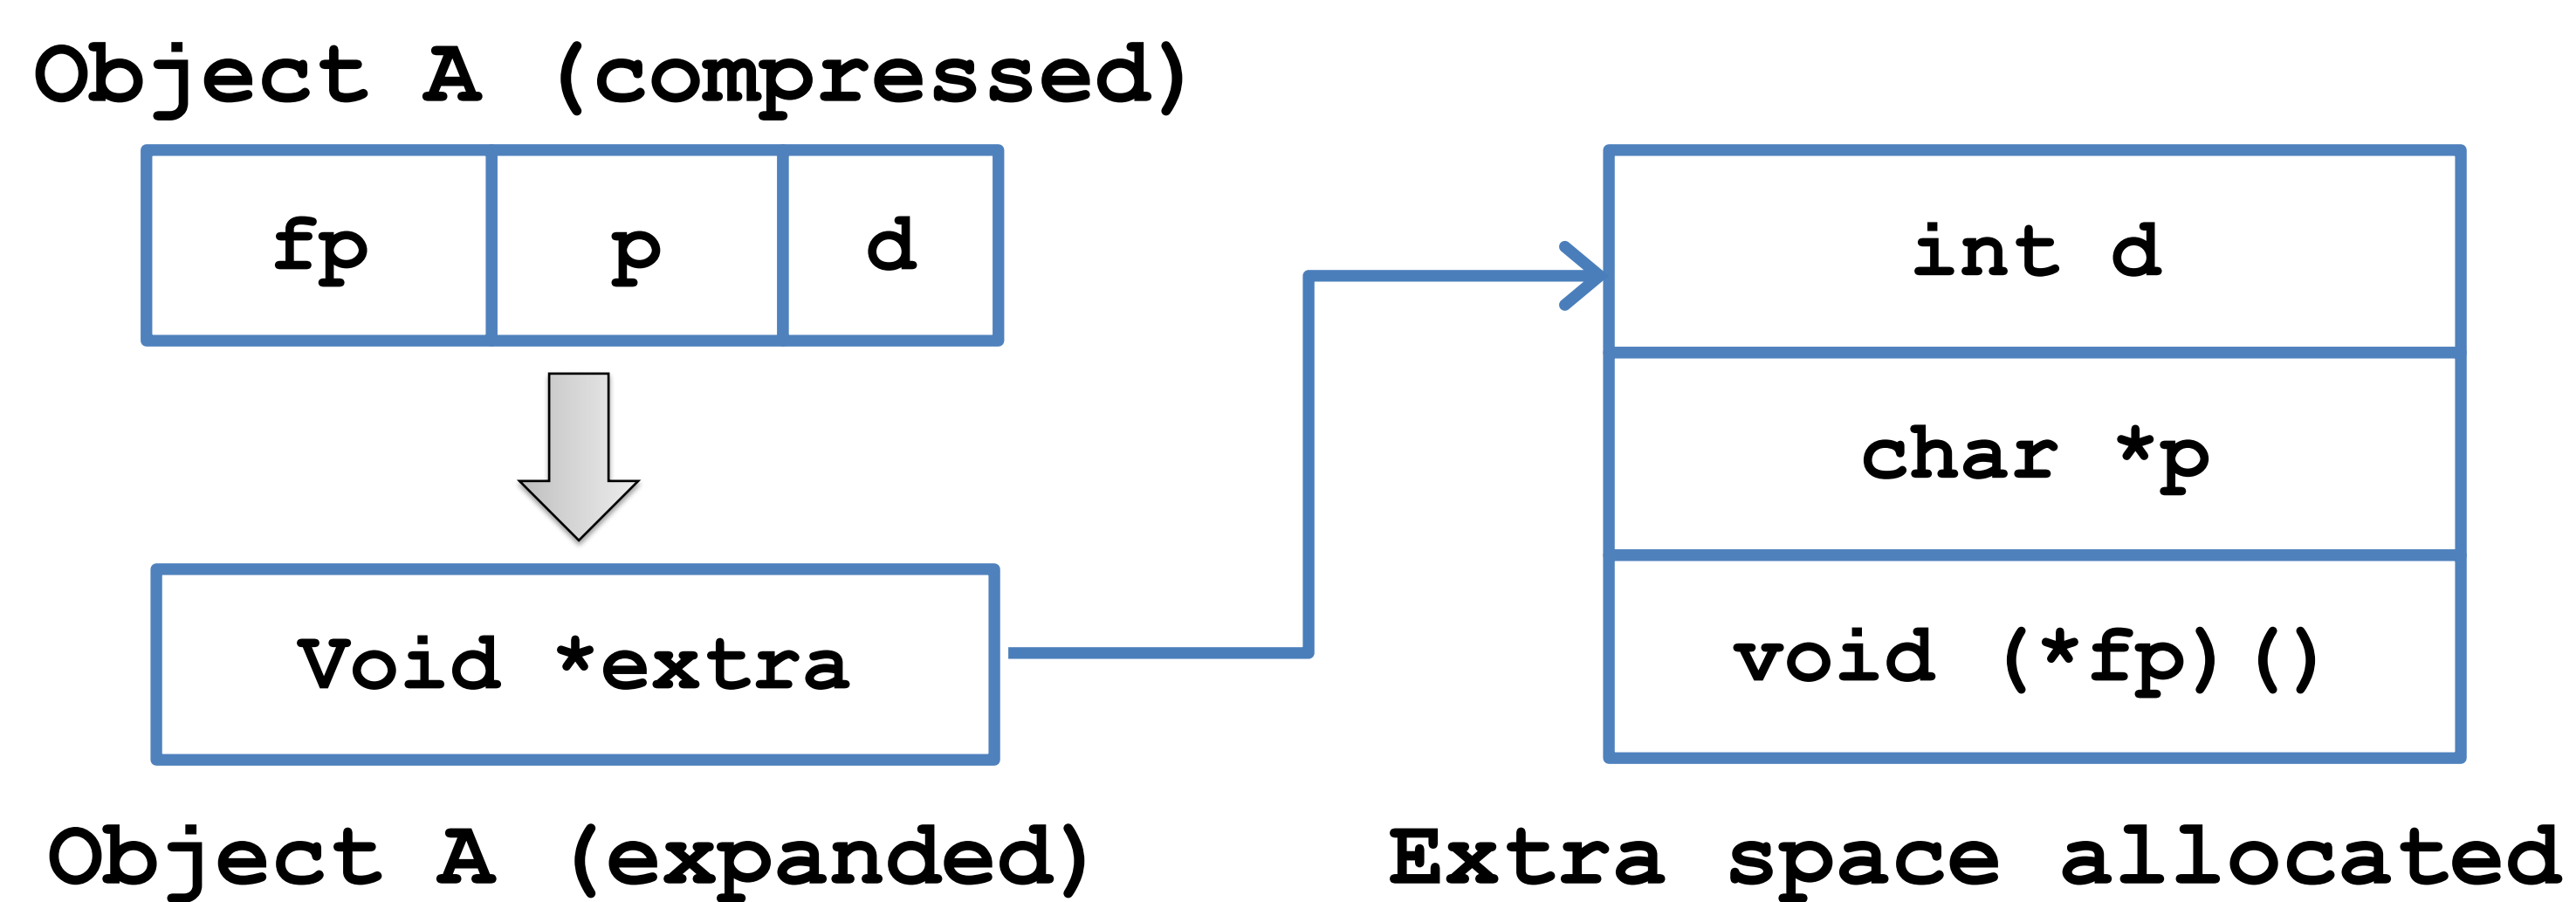


Opportunity: Most scalar variables do not require a machine word to store values.**

Compression and Decompression

- Primitives:* given a compression scheme S
- $test_compress(S, v)$, if v can be compressed with S
 - $compress(S, v)$, does the actual compression
 - $decompress(S, v)$, decompresses v with S

What if an incompressible value is encountered?
Expand compressed data object with additional memory space.

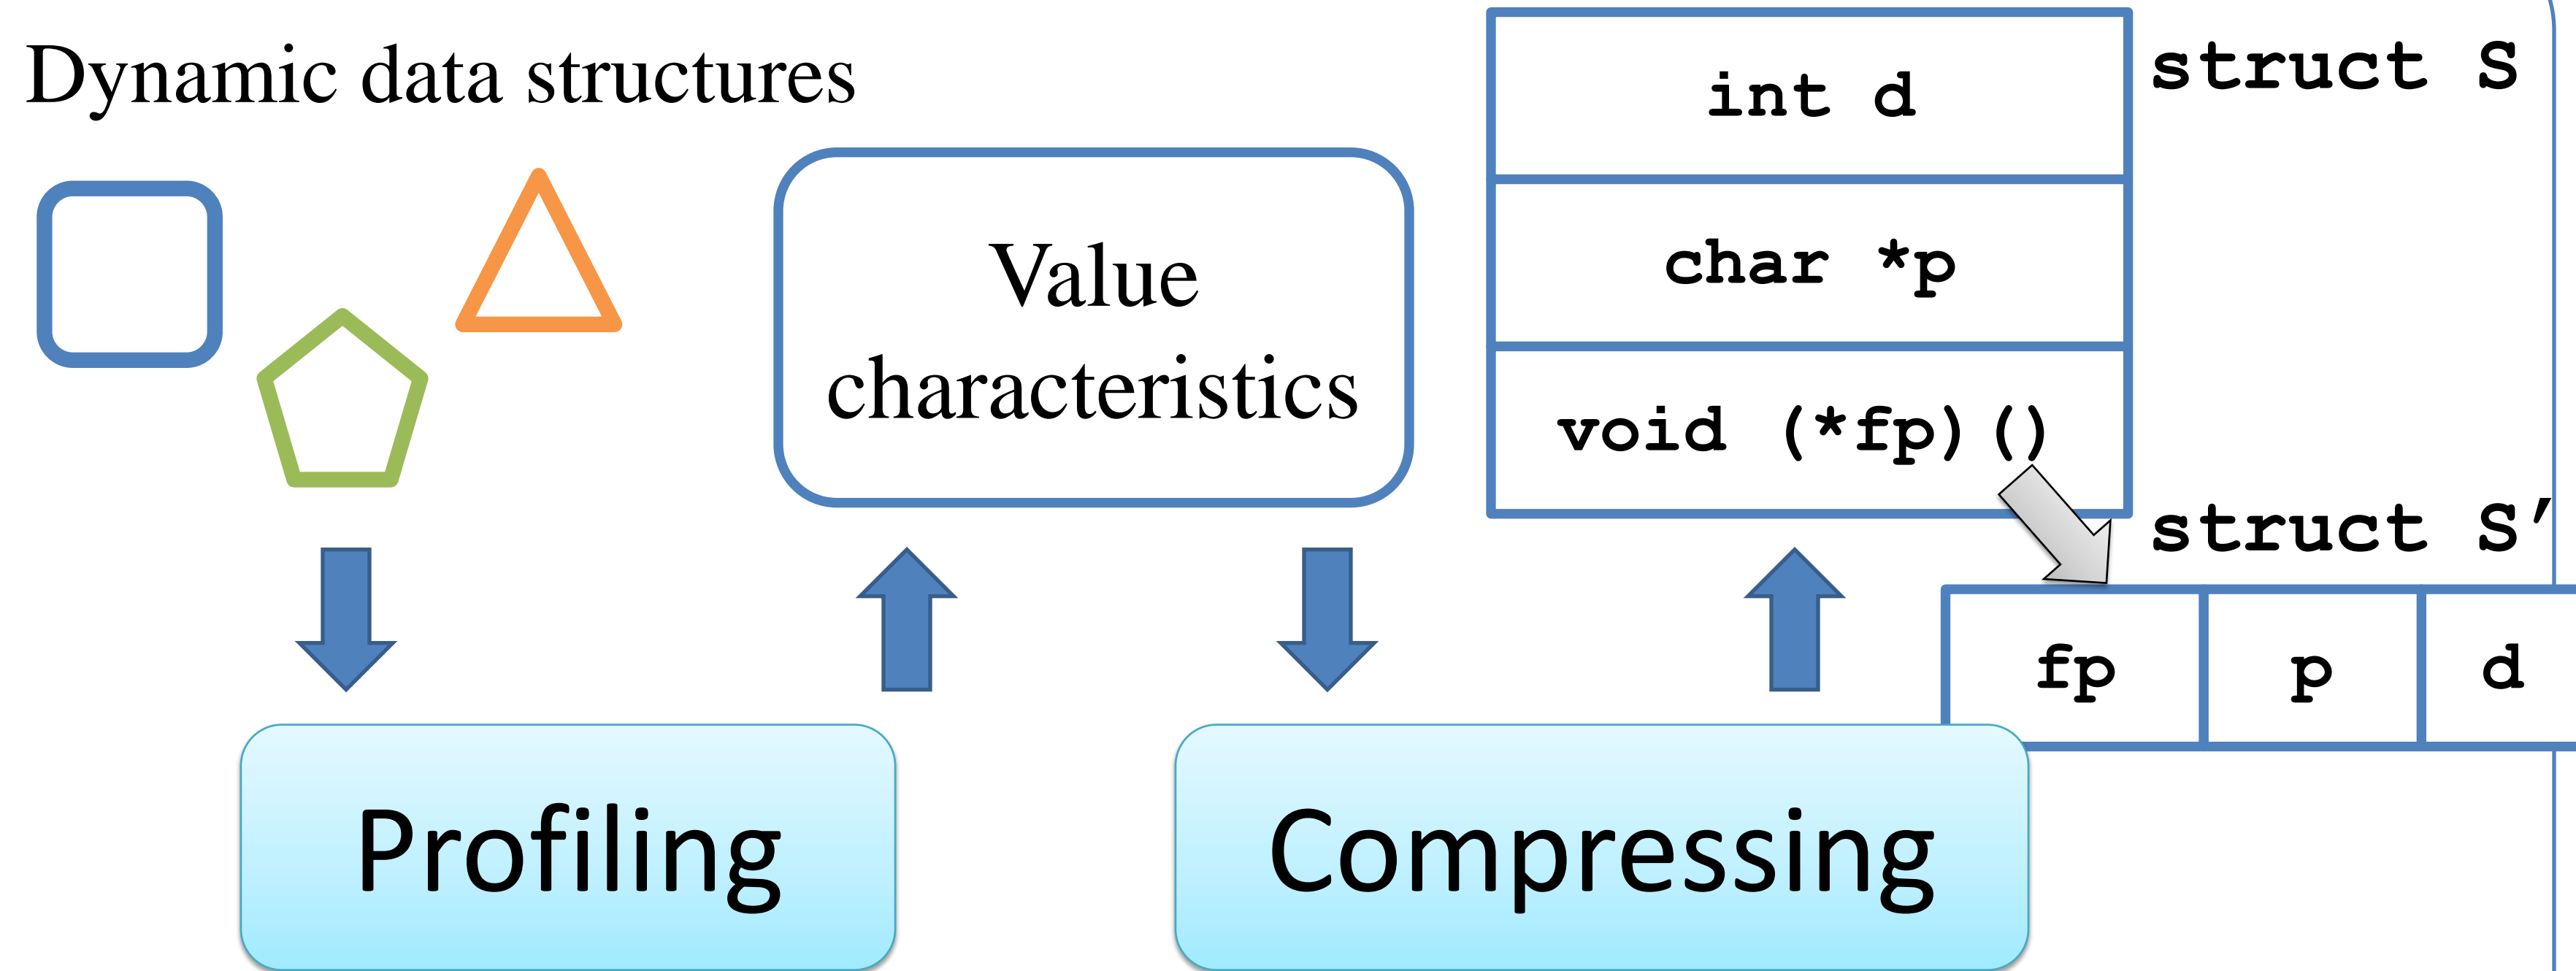


Source code transformation

Original	Transformed
<code>struct S {int d; char *p;};</code>	<code>struct S' {char cdata[size];};</code>
<code>s->d = v;</code>	<code>compress_store(s->cdata, fid, v)</code>
<code>v = s->d;</code>	<code>v = compress_load(s->cdata, fid)</code>

fid is a unique number assigned for each compressed structure field.

Data structures compression



Example

Value characteristics	Compression schemes
Narrow width [0~0x100]	9 bits
Common prefix [0x8c0000~0x8c0fff]	12 bits (offset) + prefix
Small set {1, 5, 7, 0x7fffff}	2 bits (index) + Value table

Choose the compression scheme which achieves the smallest bit width.

Experimental results



Slab allocators in Linux kernel 2.6.19

Struct Name	OrigSize	CompSize	Reduce
ext2_inode_info	372	276	25.8%
inode	264	212	19.7%
dentry	124	80	35.5%

Why these structures? These structures consumes about 80% of the total size of memory spaces in the slab allocators.

Dynamic impact Reduces the total memory consumption of the slab allocators by 16% when running *MediaBench* suite with minimal increases in execution time (1.6%).

Related work

- Ali-Reza et al.: **Improving 64-bit java IPF performance by compressing** – compressing 64-bit pointers.
- Zhang and Gupta: **Data compression transformations for dynamically allocated data structures** – hardware assisted approach.
- Coopriker and Regehr: **Offline compression for on-chip RAM** – compressing static allocated objects.

* This work was supported in part by NSF Grants CNS-0410918 and CNS-0615347.

** Brooks and Martonosi showed that over half of the integer operations in SPECint95 can be represented with 16 bits or less.