

Evaluating the Importance of User-Specific Profiling

Zheng Wang

Harvard University

Norm Rubin


Digital Equipment Corporation

Simply Put ...

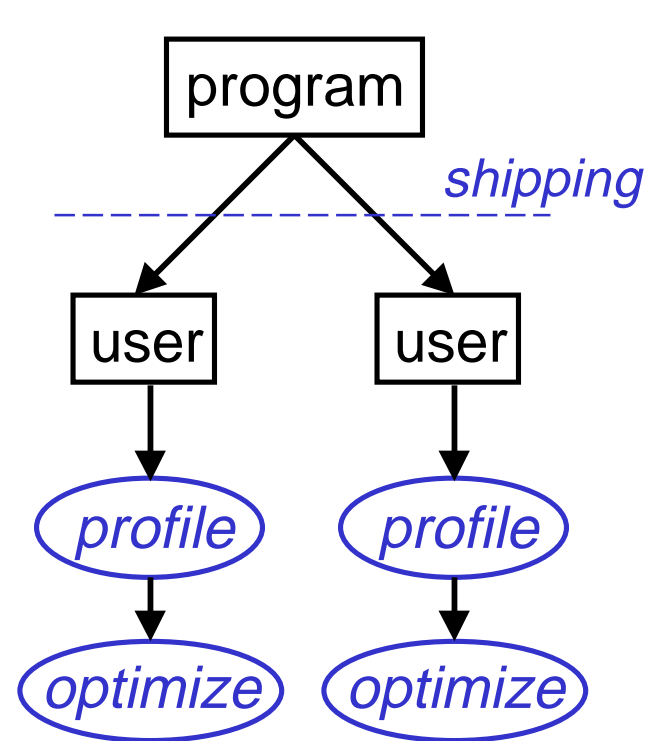
How
different users
use the
same program

Background

Two models for applying profile-based optimization



Assumption:
users are similar



Assumption:
users are different

Our Goal

- Target interactive applications on Windows NT
- Compare the *usage patterns* of different users
- Examine the impact of differences in profiles on the optimization performance

FX!32

- DIGITAL FX!32 emulator / binary translator
 - automatically runs x86 applications on Alpha NT
- FX!32 profiles
 - accumulated over multiple runs
 - contain information on
 - procedure calls
 - indirect control transfers
 - unaligned memory references

Benchmarks

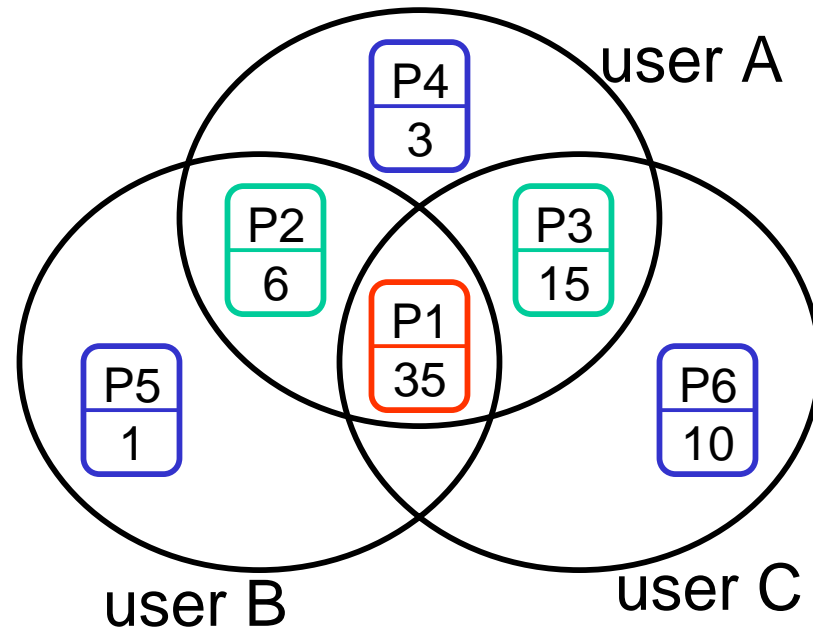
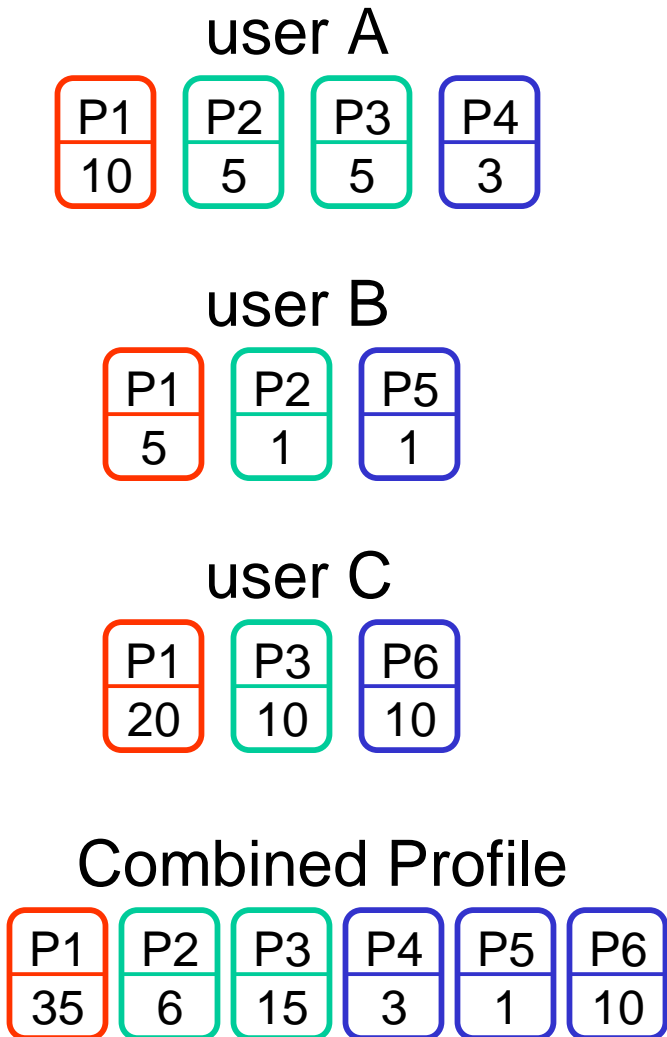
Interactive application executables / DLLs

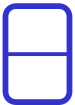
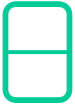
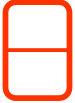
Benchmark Module		Description	File Size (KB)
Microsoft Office	excel.exe (97)	Office 97 Excel main executable	5469
	mso95.dll (95)	Office 95 (Version 7.0) DLL	918
	mso97.dll (97)	Office 97 DLL	3686
	outllib.dll (97)	Office 97 Microsoft Outlook DLL	4254
	powerpnt.exe (97)	Office 97 PowerPoint executable	3411
	winword.exe (95)	Office 95 Word executable	3755
	winword.exe (97)	Office 97 Word executable	5194
Others	acrord32.exe	Adobe Acrobat Reader 3.0 executable	2265
	mfc40.dll	Microsoft Visual C++ 4.0 DLL	901
	netscape.exe	Netscape Navigator Gold 3.01 executable	3093
	photoshp.exe	Adobe PhotoShop 4.0 executable	3560
	pnui3250.dll	Support library for RealPlayer (32-bit) 5.0	590
	winhlp32.exe	Windows NT 4.0 help utility	303
	winzip32.exe	WinZip compression utility 6.2	846

Methodology

- Profile collection
 - generated from users' spontaneous usage
 - over 20 users in total
 - 4-12 users for each benchmark module
 - time span of several months
- Statistical analysis
 - compare the individual profiles for each module
 - focus on the set of procedures each person uses

Profile Comparison

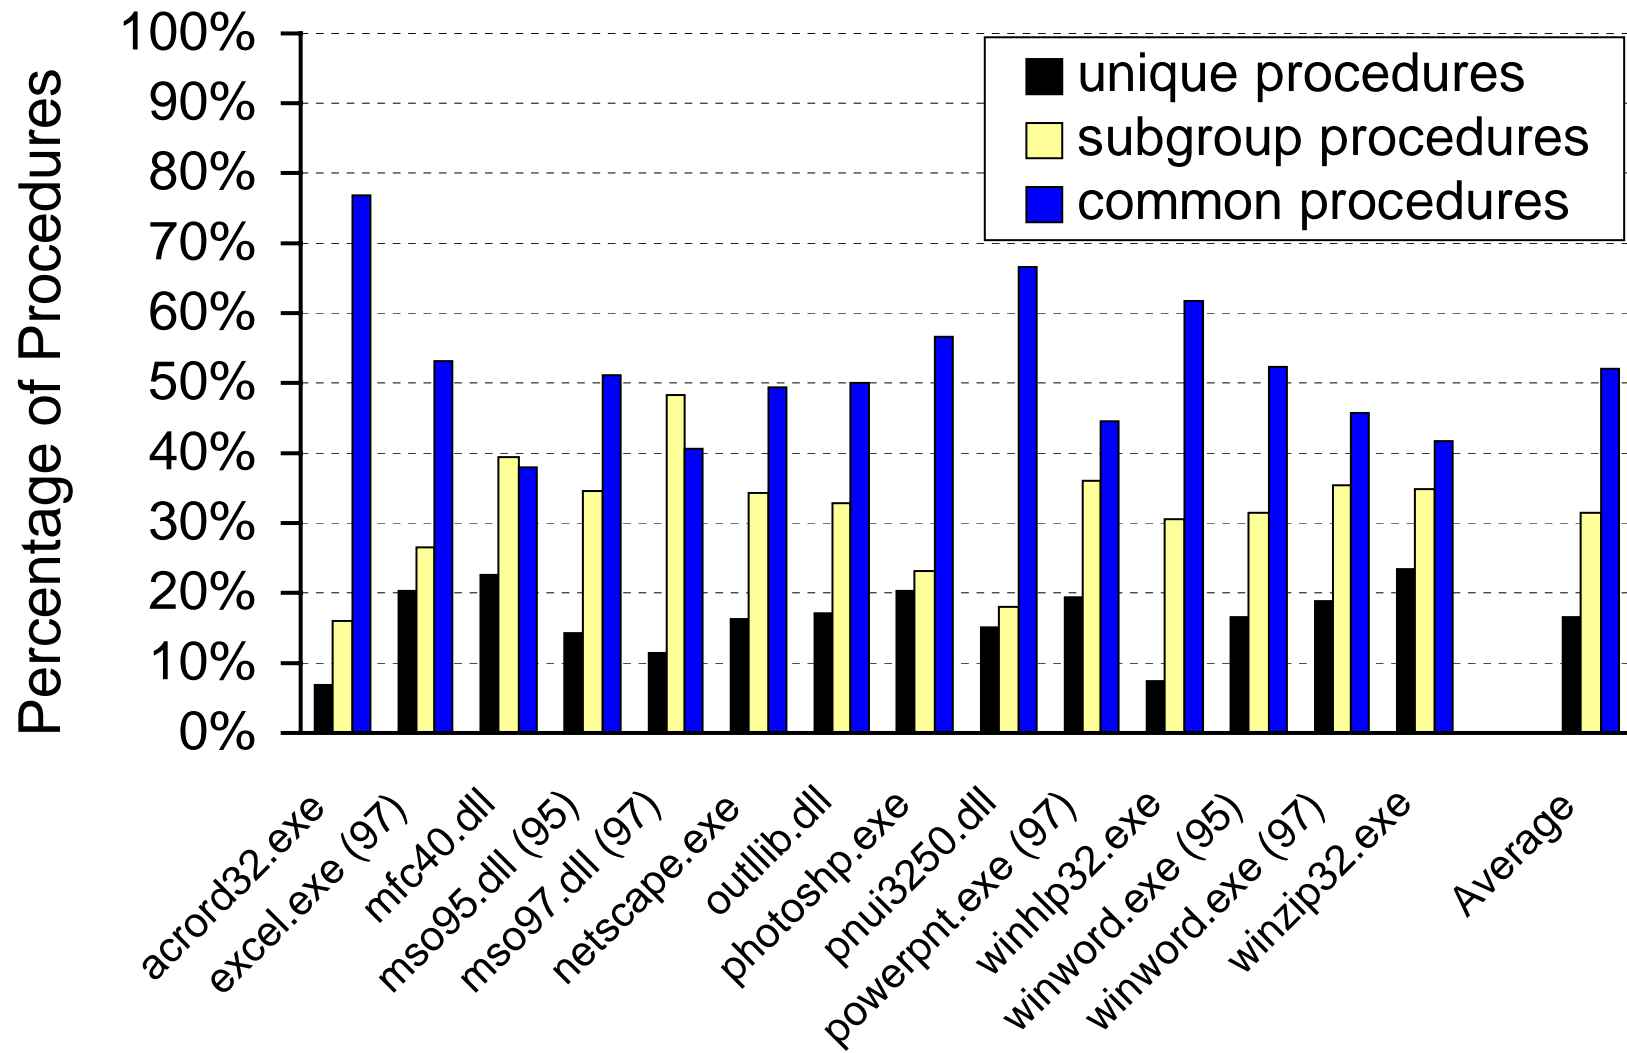


-  unique procedure
(usage count = 1)
-  subgroup procedure
-  common procedure
(usage count = number of users)

Summary of Profiles

Benchmark Module	Number of Users	Number of Procedures in <i>Combined</i>	Percentage of Procedures in an Individual Profile		
			Smallest	Largest	Average
acrord32.exe	4	5050	79.4%	94.9%	87.8%
excel.exe (97)	4	8514	69.1%	80.1%	74.6%
mfc40.dll	7	2558	49.3%	75.5%	60.2%
mso95.dll (95)	6	2630	64.4%	80.4%	73.3%
mso97.dll (97)	8	9994	56.3%	83.0%	68.7%
netscape.exe	4	7938	61.1%	88.2%	73.7%
outllib.dll (97)	5	16330	64.6%	80.3%	72.9%
photoshp.exe	5	10502	66.5%	84.2%	74.3%
pnui3250.dll	4	1443	73.1%	90.0%	81.8%
powerpnt.exe (97)	5	15014	59.3%	85.9%	70.0%
winhlp32.exe	12	762	71.3%	95.4%	79.0%
winword.exe (95)	5	7317	62.9%	85.0%	73.8%
winword.exe (97)	6	10113	61.6%	79.0%	68.9%
winzip32.exe	5	1125	53.1%	79.4%	65.5%
		Average	63.7%	84.4%	73.2%

Procedure Distribution



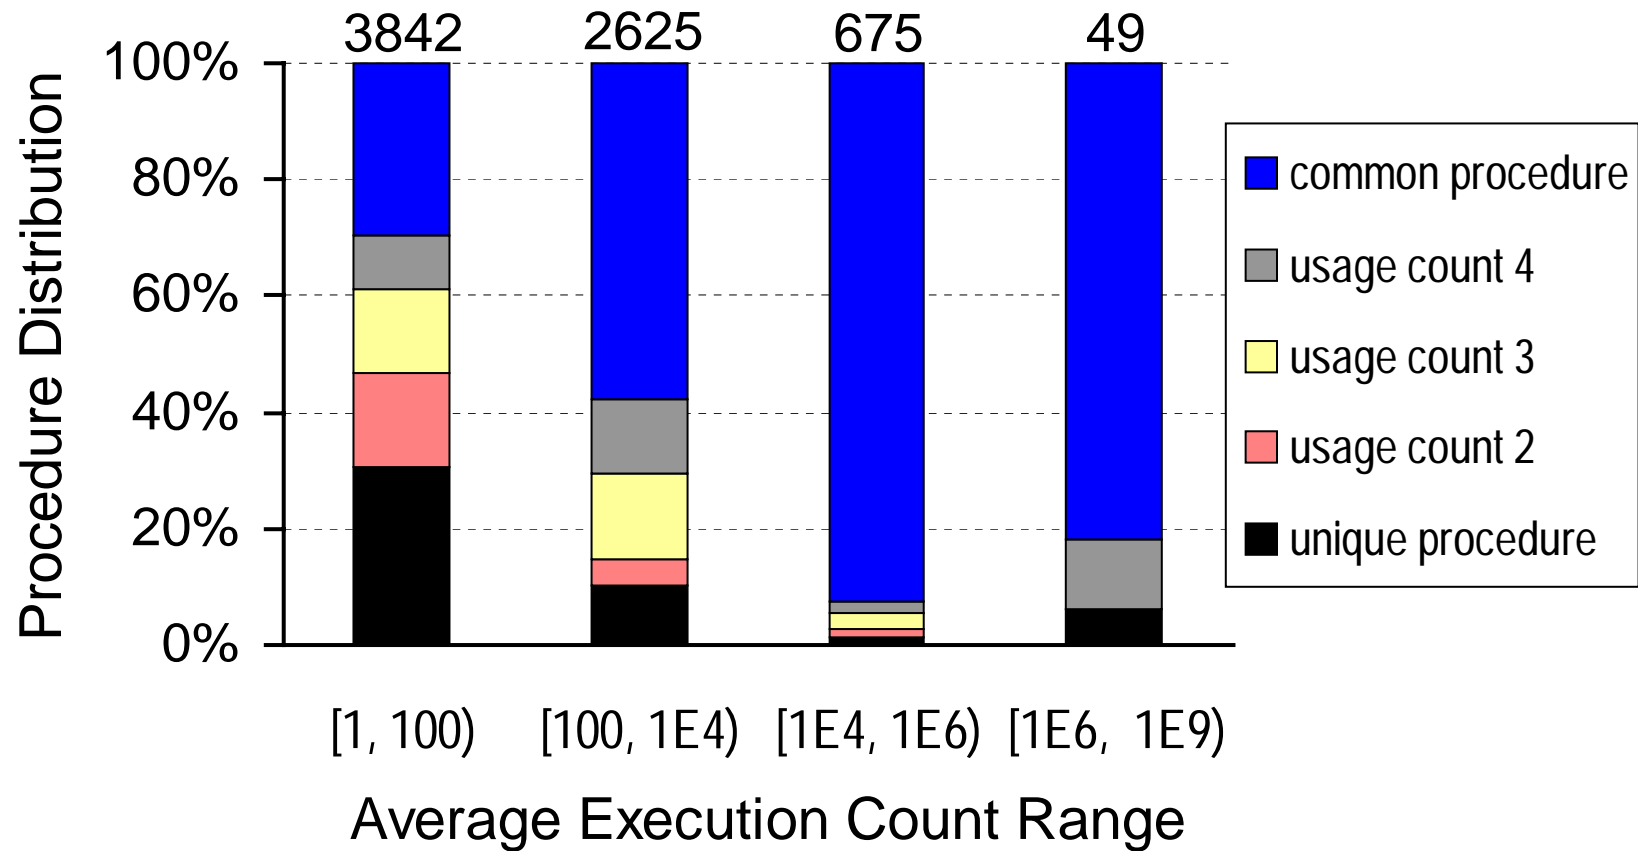
Pair-Wise Profile Comparison

	Happy	Sneezy	Grumpy	Doc	Bashful
Bashful	66.6%	69.1%	69.6%	77.2%	--
Doc	72.5%	69.7%	73.5%	--	
Grumpy	71.6%	76.4%	--		
Sneezy	76.0%	--			
Happy	--				

winword.exe (95)


$$\text{Similarity \%} = \frac{\text{number of procedures included in both profiles}}{\text{number of procedures included in either profile}}$$

Usage Count vs. Execution Count



When Profiles Grow

Date	Number of Procedures Used By						Procedure Distribution		
	Bashful	Doc	Grumpy	Sneezy	Happy	<i>Combined</i>	Unique	Subgroup	Common
10/10	4600	4091	4691	5648	6222	7191	20.4%	33.5%	46.2%
10/15	4600	4465	4691	5648	6222	7213	19.5%	33.2%	47.3%
10/22	4600	4465	4947	5648	6222	7239	18.2%	34.0%	47.8%
10/29	4600	4834	5332	5648	6222	7283	17.0%	31.6%	51.8%
11/03	4600	4990	5332	5648	6222	7288	16.3%	31.5%	52.3%
11/10	4600	4990	5332	5846	6222	7317	16.5%	31.2%	52.2%



Change in similarity over time

winword.exe (95)

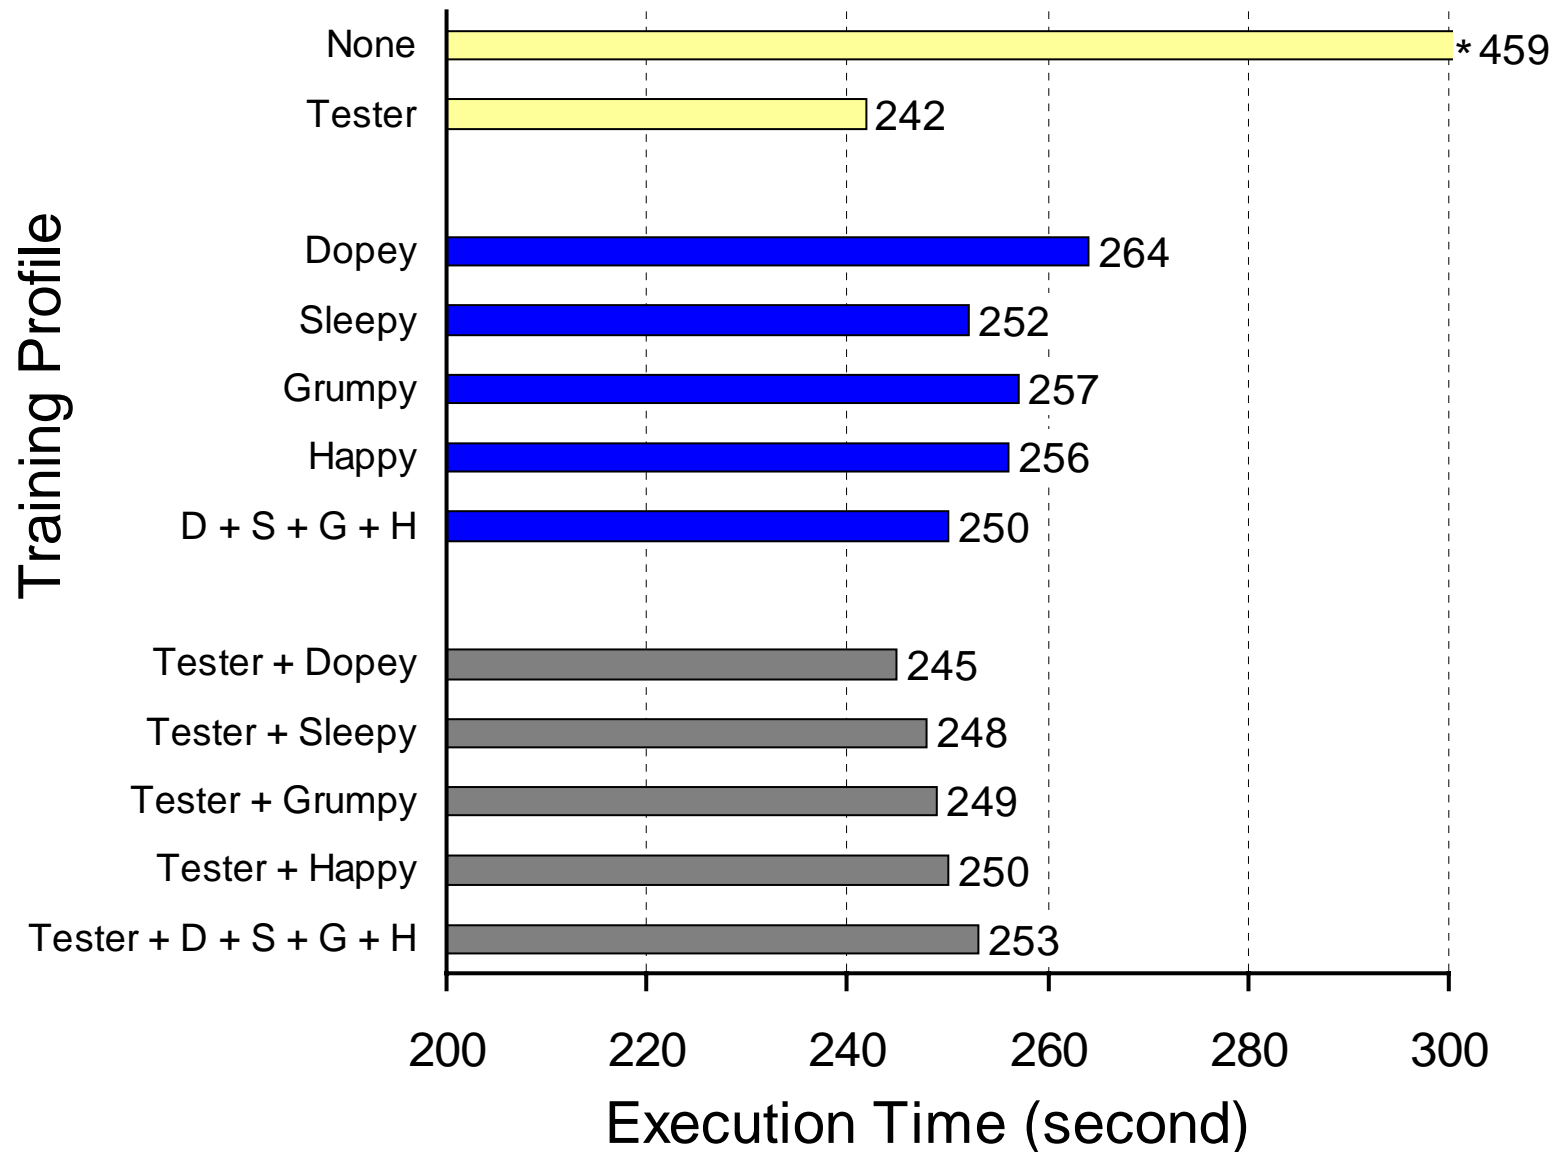
Optimization Performance

Question: How much impact the differences in profiles have on optimization performance?

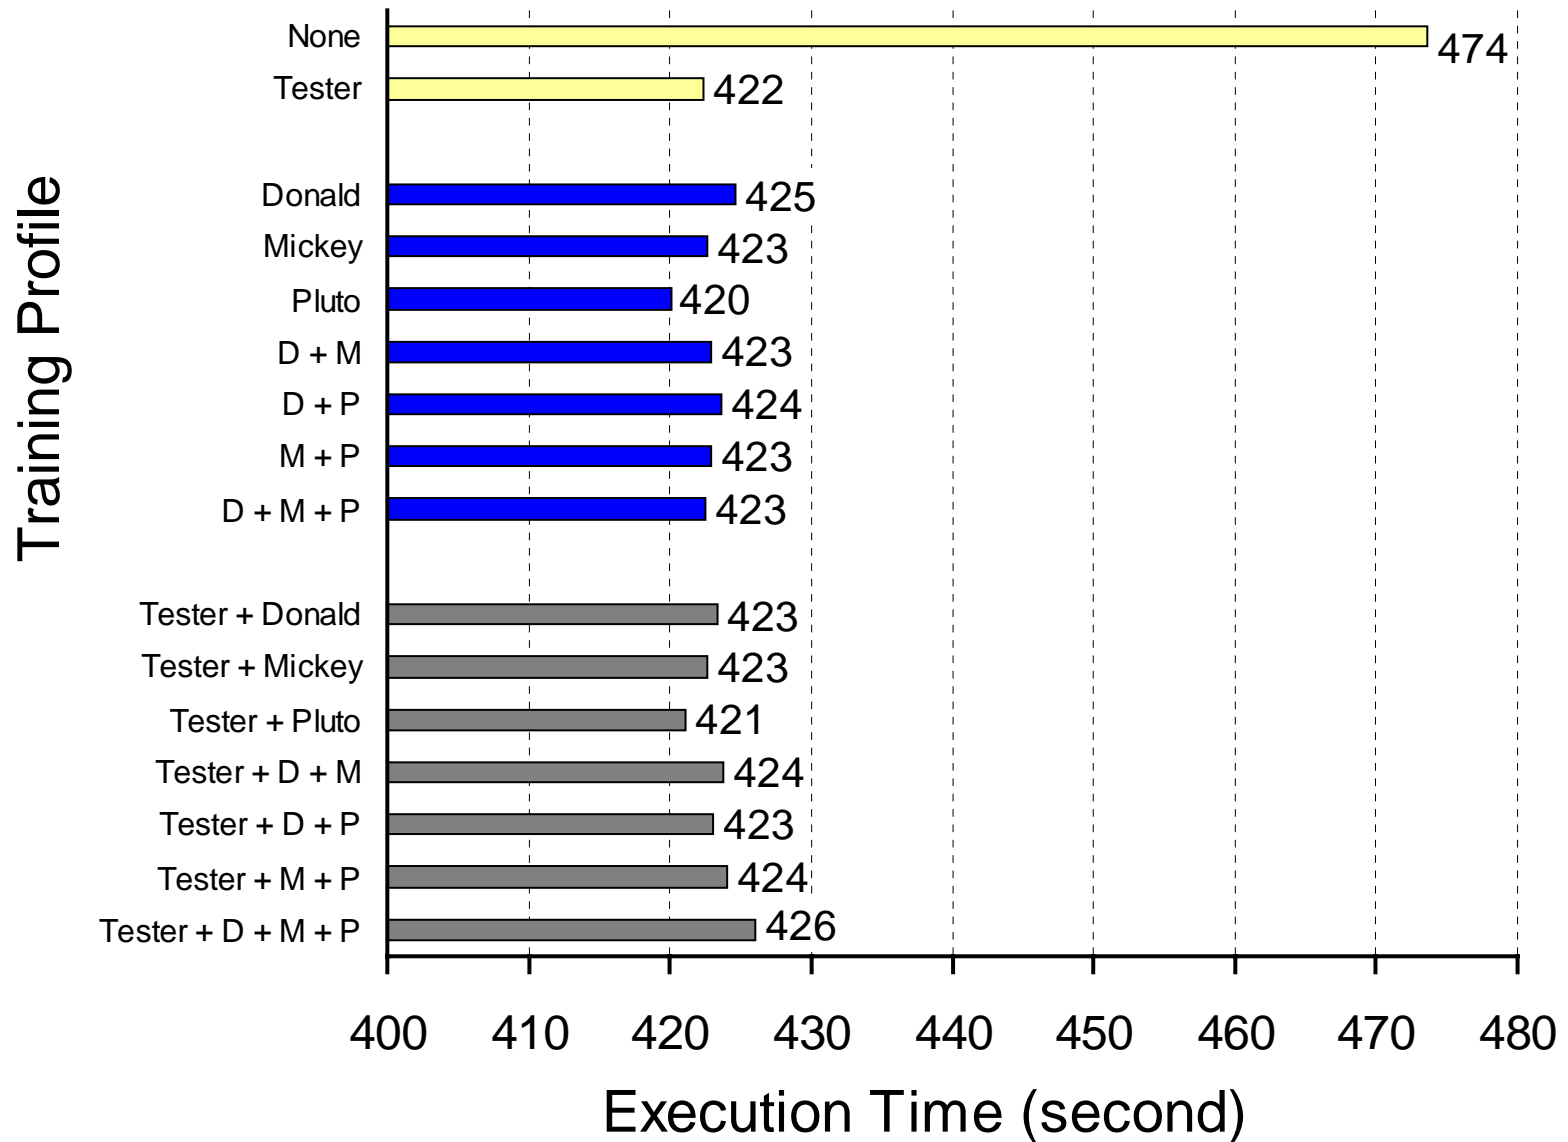
- Our case: FX!32 program translation/optimization
- Different Individual and group profiles for training
- Automated test script for performance measurement

- Two benchmarks show different results
 - winword.exe (95)
 - powerpnt.exe (97)

Optimization Results: Word 95



Optimization Results: PowerPoint 97



Summary

- People use different sets of procedures in a program
 - on average, 52% of used procedures are *common* and 16% are *unique*
- Frequently executed procedures tend to be common
- Some differences persist over time

- Differences in profiles can have measurable impact on optimization performance
- Profiles from another user or a group can be less effective