CE 817 - Advanced Network Security
VoIP Security

Lecture 25

Mehdi Kharrazi
Department of Computer Engineering
Sharif University of Technology

Acknowledgments: Some of the slides are fully or partially obtained from other sources. Reference is noted on the bottom of each slide, when the content is fully obtained from another source. Otherwise a full list of references is provided on the last slide.

Background: the PSTN

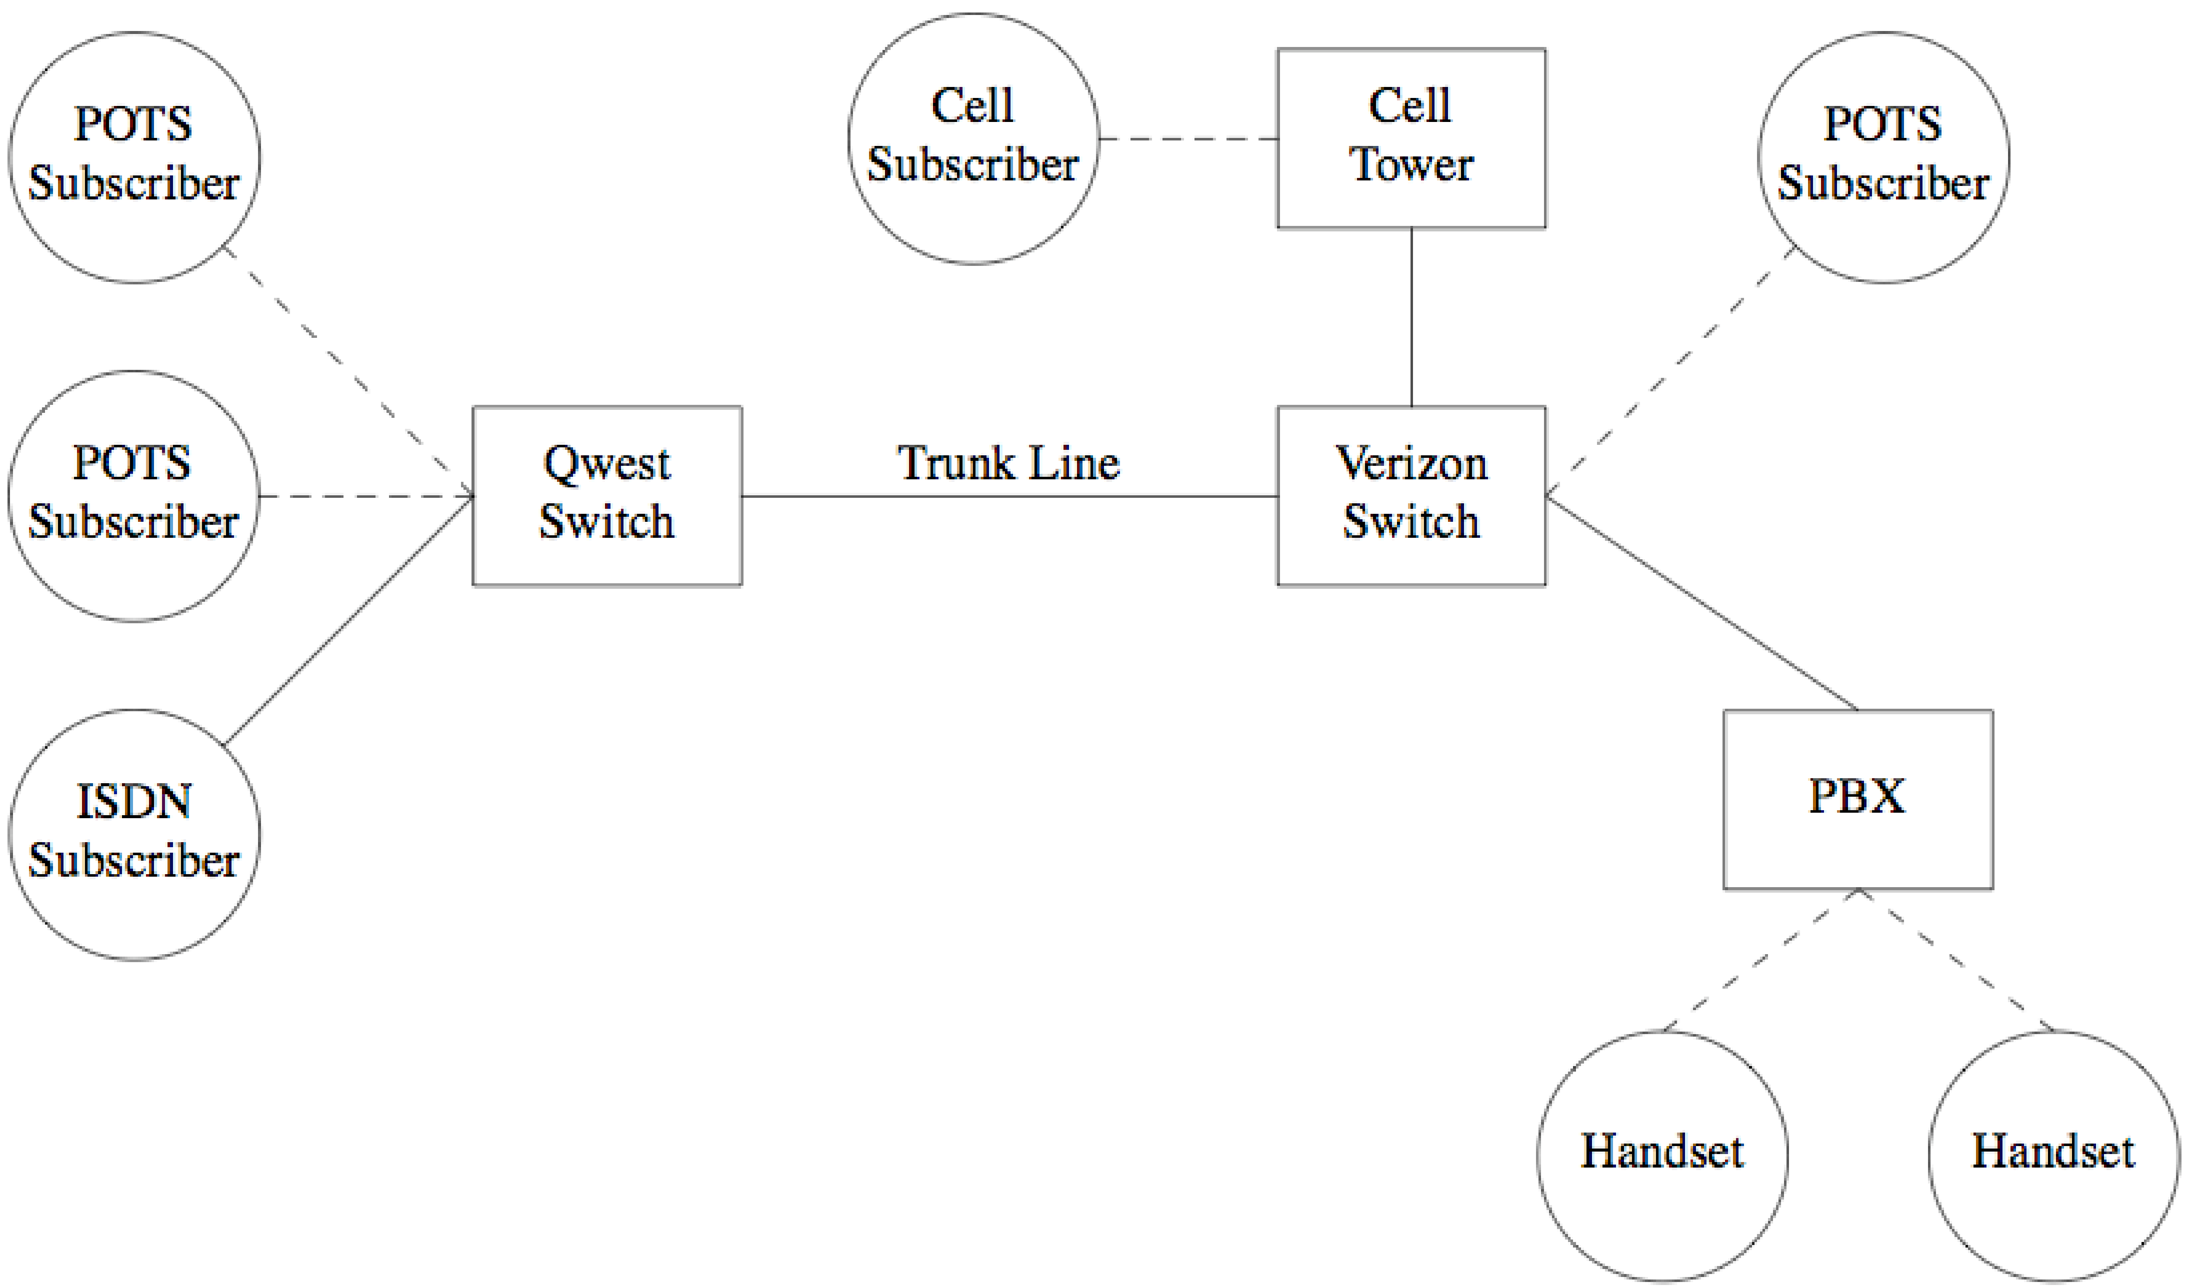

Plain Old Telephone Service(POTS)

- This is what you probably have
- Analog transmission
  - A pair of copper wires from you to the CO
- All signalling is inband
  - Instructions from you to the switch are DTMF tones
  - From the switch to you is tones (e.g., caller ID)
- Basically no security
  - Wiretapping means a pair of alligator clips and a speaker
  - Hijacking is just as easy

What is SIP?

- Session Initiation Protocol
- Control channel for Voice over IP
- (Other control channel protcols exist, notably H.323 and Skype's, but we'll focus on SIP)

What’s a Control Channel?

- A control channel — known in the telephone world as a signaling channel — does call setup
- It locates the other end point, determines if it’s available, asks the endpoint to alert the called party, passes back status to the caller, etc.
- Even in a pure IP world, we need a signaling channel; when connecting to the PSTN (Public Switched Telephone Network), it’s essential

History of Signaling Channels

- Telephone signaling was once done “in-band” — that is, the pulses or tones were sent over the same circuit as would later be used to carry the voice traffic for that call
- “Blue boxes” — telephone fraud devices — worked by simulating some of the control tones used to set up free calls
- The solution was to move signaling to a separate, “out-of-band” data network, known today as CCIS (Common Channel Interoffice Signaling)
- Out-of-band signaling is more efficient; it allows easy creation of fancier services

Signaling and VoIP

- Why can't we just call a domain name or IP address?
- Example: Many endpoints don't have stable, easily-memorized domain names
  - IP addresses change frequently, especially for dial-up and hotspot users

Complexity

- PSTN interconnection: very many endpoints have just a few IP addresses
- Besides, someone has to pay for the PSTN interconnection
- Firewalls
- Network address translators (NATs)
- Mapping between “phone number” and IP address
- Business arrangements between telephone companies
- Unreachable hosts
- Fancy phone features

[Bellovin06]

Basic SIP Architecture

- SIP endpoints speak IP
- Ideally, the actual conversation would be end-to-end, from one SIP phone to the other
- Each node can use a SIP proxy for call setup

Simple SIP Calling

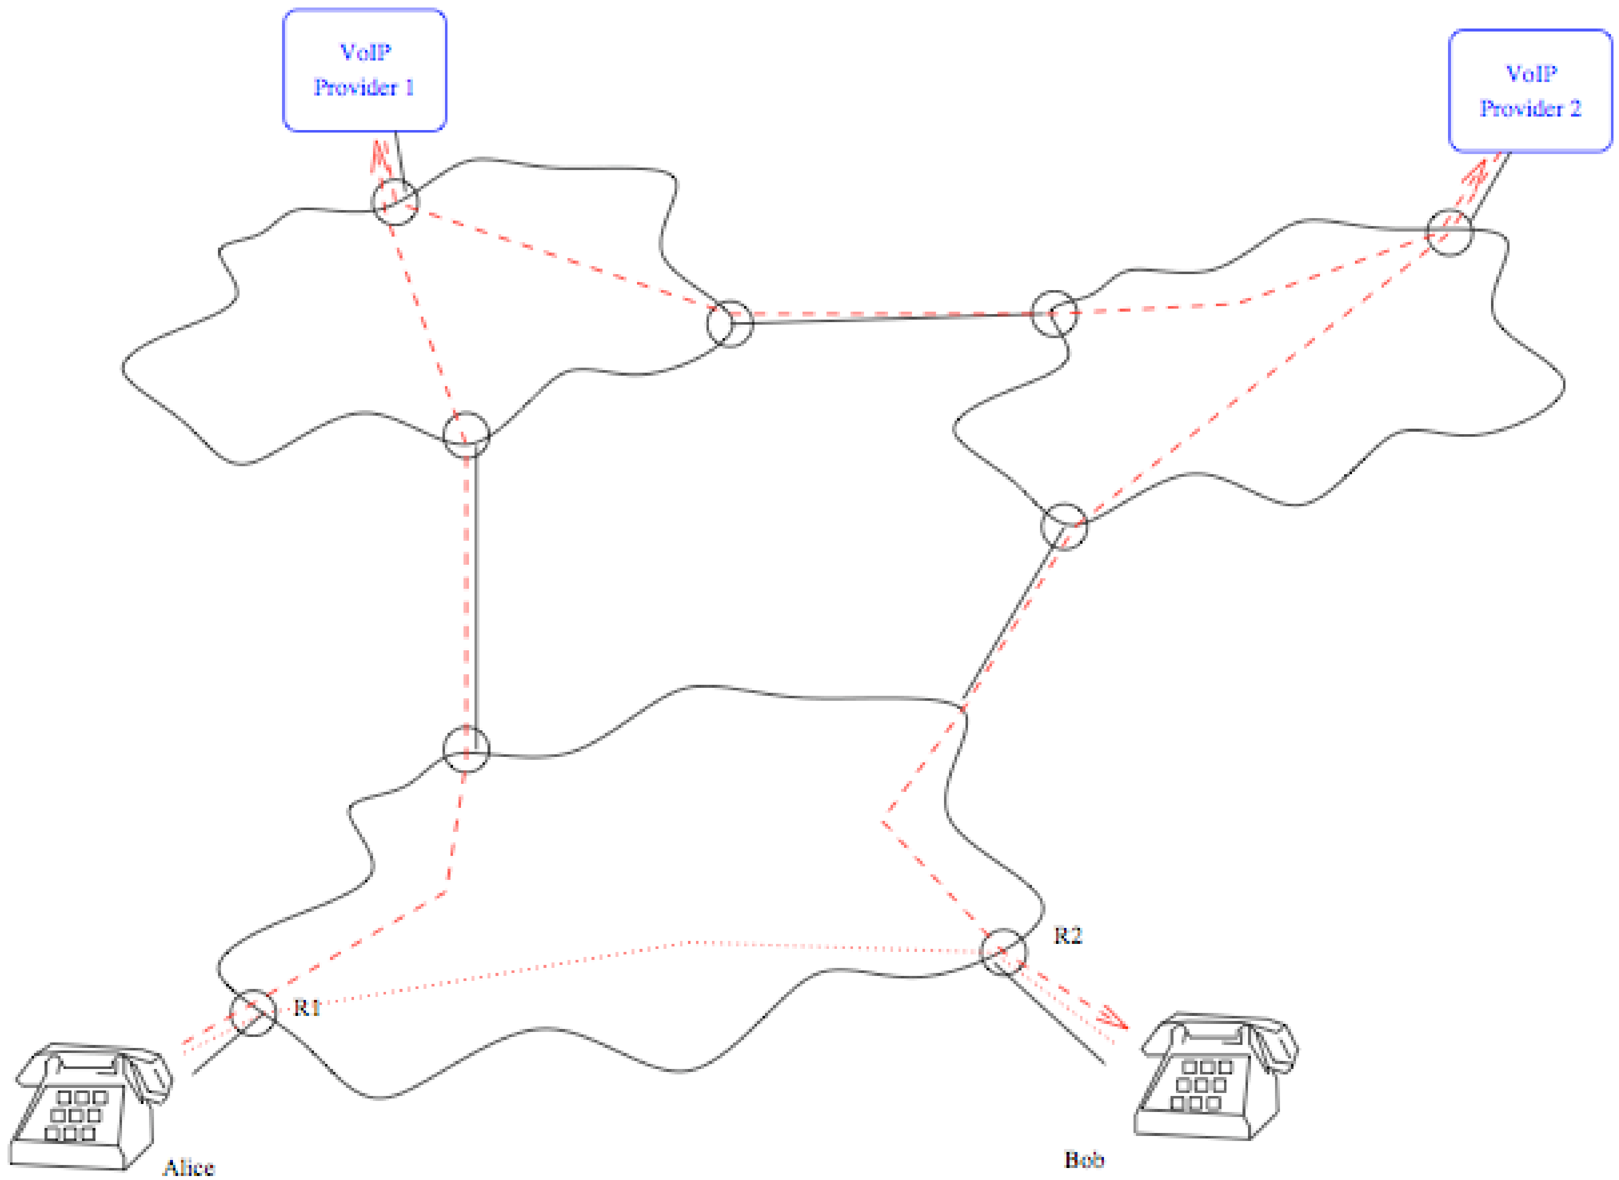

Alice Calls Bob

- Alice uses VoIP Provider 1 (VP1) as her proxy; Bob uses VoIP Provider 2 (VP2) as his
- To call Bob, Alice sends a SIP URI to VP1 via TCP
- VP1 determines that the URI points to VP2, so the calls setup request is relayed there via TCP
- VP2 tells Bob about the call via TCP; if he wants to, he can accept it
- Notification is sent back to Alice via VP1
- Alice establishes a direct UDP data connection to Bob for the voice traffic

SIP URIs

- How is a SIP URI converted to a SIP proxy address?
- What about ordinary telephone numbers?
  - tel: URIs are used for ordinary phone numbers
- Example SIP URI :
  - SIP: someone@example.com
- Example tel: URI :
  - TEL: + 0 216 - 616 - 4601
- All SIP URIs are converted by means of DNS magic: NAPTR records
- (For this class, the details aren't important — the essential point is that by means of repeated, complex DNS lookups, any SIP URI is converted to an IP address)

Attacking SIP

The Usual Questions

- What are we trying to protect?
- Against whom?

Information at Risk

- Voice content itself
- Caller and called party for each connection
- Billing information

[Bellovin06]

Voice Content

- Confidentiality is the main concern
- Is VoIP easier to wiretap than traditional phone service?
- Only the endpoints should see that information; can be encrypted through proxies
- Relatively hard to spoof a voice in real-time, so authenticity is not a major concern

 [Bellovin06]

Caller/Called Party Information

- Of great interest to many parties (look at the HP case — that's the data HP was after)
- Useful even after the call (you can't intercept a call after it's over; you can look at who talked)
- Must be kept confidential — but proxies need to see it, to route the call
- Must be authentic, or the call could be misrouted maliciously

Billing Information

- Derived in part from caller/called party information
- May have other information from call routing process
- As before, must be confidential
- Integrity failures can lead to billing errors, in either direction
- (Often a major privacy concern after the fact — again, consider the HP case.)

Eavesdropping on a Call

- Simplest approach: listen on some link
- Which link is best for targeting a given person?
- Easiest: their access link
- What if they're mobile? Hard — they could be coming from anywhere
- Do you have the physical ability to listen on the VoIP provider's links? What if the VoIP provider is in a distant, unfriendly country?

Registration Hijacking

- An attacker can try to register with VP2 as Bob
- If the attacker succeeds, all calls destined for Bob with be routed to the attacker

Abusing the DNS

- Call routing is partially controlled by the DNS
- Is it possible to corrupt the DNS answers?
- By creating fake DNS entries, it's possible to reroute the call to go via an intercept station

[Bellovin06]

Caller/Called Party Information

- Again, link eavesdropping and DNS attacks are straightforward
- The task is easier here; proxies (usually) don't move around
- VoIP providers are high-value targets, since they process many calls

Hacking the Proxies

- Is it possible to hack the VoIP proxy servers?
- Sure — why not?
- Conventional phone switches can be (and some are) hacked, but there's a big difference: the attacker can speak a much more complex protocol to a SIP switch than to a PSTN switch, which means they're more vulnerable
- It's hard to do too much damage with just a few touch-tones!
- Aside: fancier services are easier to hack, on both kinds of telephone systems

Defenses

Protecting SIP

- As usual, we'll use crypto to guard against eavesdropping
- The details, though, are tricky

Alice to VP1

- Alice has a trust relationship with her proxy
- Authentication is relatively easy
- Usually, TLS is used to protect the TCP session to the proxy
- Alice must verify VP1's certificate
- Alice can use passwords or client-side certificates to authenticate herself

[Bellovin06]

Proxy to Proxy Traffic

- VP1 may not have a trust relationship with VP2
- How can VP1 get VP2's certificate?
- More precisely, how can VP1 validate it, if they don't share a trust anchor?
- This applies regardless of what security protocol is used (though TLS is the norm)

End-to-End Signaling Traffic

- Some signaling traffic must be secure end-to-end
- Example: Bob needs to know, authoritatively, that it's Alice who has called him
- However, the intermediate nodes need to see this
- Solution: digitally sign the data (using S/MIME), but don't encrypt it

 [Bellovin06]

Key Management for the Voice Call

- How do Alice and Bob get a shared key for voice traffic encryption?
- Alice uses S/MIME to send Bob an encrypted traffic key
- But — how does Alice get Bob's certificate?
- There is no general PKI for SIP users
- True end-to-end confidentiality can only happen by prearrangement
- (This statement is more generally true. . .)

The State of Practice

- Most vendors don't implement the fancy crypto
- VoIP is thus not as secure as it could be (but Skype does do a lot of crypto)
- NIST recommends great care in using VoIP — see http://csrc.nist.gov/publications/nistpubs/800-58/SP800-58-final.pdf

Caller ID

CallerID

- Suppose the SIP call is being relayed to the PSTN
- Where does the CallerID information come from?
- Can it be spoofed?

Phone Network Design

- The phone network was based on trust — only “real” telephone companies had phone switches
- No authentication was done on information from other switches, including CallerID
- Today, anyone can run a phone switch. . .

[Bellovin06]

CallerID and VoIP

- Run Asterisk, an open source PBX program, on some machine
- Get a leased line to a VoIP-to-PSTN gateway company
- Configure Asterisk to send whatever information you want. . .
- This abuse is happening now

SPIT (Spam Over IP Telephony)

Background

- SPAM considered one of biggest problems in Internet
- SPIT is expected to become a major issue in the next few years with increasing deployment of VoIP solutions
- Potential for productivity disturbance is much greater than SPAM

[Santos]

Background

- Definition: The transmission of unsolicited calls over Internet telephony (VoIP)
- "SPITTERS" will forge their identities
- SPITTING agent capable of placing hundreds of simultaneous automated calls

SPAM vs. SPIT

SPAM	SPIT
User can sort through or filter messages based on content and header	VoIP is a real time protocol that does not allow grant the receiver access to the contents of the call prior to its acceptance
Email is delivered asynchronously, whenever a user decides to download/access email	Victim is interrupted instantly with the phone ringing
SPAMMER does not know for sure when or whether his message will reach the victim	A successful call guarantees that the user exists, is currently online, and will most likely receive the message soon.

[Santos]

SPIT Prevention Framework

- Goals:
  - Minimize false positives & negatives
  - Minimize callee interaction in identifying SPIT
  - Minimize inconvenience to caller
  - General enough to work in different environments (work, home, etc) and cultures

SPIT Prevention Framework

- 5 Stage Approach:
  - Stage 1: no interaction w/ users
    - Blacklist, Whitelist, Graylisting, Circles of Trust, Pattern / Anomaly Detection
  - Stage 2: caller interaction
    - Computational Puzzles, Sender Checks, Audio CAPTCHAS

[Santos]

SPIT Prevention Framework

- 5 Stage Approach (continued):
  - Stage 3: feedback before call
    - Manual authorization to receive call and/or authenticate user
  - Stage 4: during the call
    - Content analysis (not currently viable)
  - Stage 5: feedback after call
    - Ex: Require a refundable payment for each call from an unknown party. The payment is only refunded if the caller was not a SPITTER.

The END

What did we cover/learn in this semester?

Lectures

- Threats and Attacks
- Firewalls
- IDS
- DoS
- Worms
- Botnets
- Honeypots
- Spyware
- Phishing

Lectures (con't)

- Traffic Analysis
- Anonymity
- Routing Security
- Network Forensics
- Wireless Security
- VoIP Security

Acknowledgments/References

- [Bellovin06] COMS W4180 — Network Security Class Columbia University, Steven Bellovin, 2006.
- [Santos] SPAM OVER IP TELEPHONY (SPIT). Identification and prevention Techniques ECE 4112 – Internetwork Security, Felipe Santos, Manoj Deshpande, Georgia Institute of Technology.